

AD-A175 784


LEARNING PHYSICAL DOMAINS: TOWARD A THEORETICAL
FRAMEWORK(U) ILLINOIS UNIV AT URBANA DEPT OF COMPUTER
SCIENCE K D FORBUS ET AL DEC 86 UTUCDCS-R-86-1247
N00014-85-K-0559


1/1

UNCLASSIFIED

F/G 5/10

NL





12

REPORT NO. UIUCDCS-R-86-1247

UIU-ENG-86-1774

LEARNING PHYSICAL DOMAINS: TOWARD A THEORETICAL FRAMEWORK


by

Kenneth D. Forbus
Dedre Gentner

December 1986

DTIC
SELECTE
JAN 07 1987

AD-A175 784



DEPARTMENT OF COMPUTER SCIENCE
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN · URBANA, ILLINOIS

ADA175784

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified		1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY		3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; Distribution unlimited	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE		5. MONITORING ORGANIZATION REPORT NUMBER(S)	
4. PERFORMING ORGANIZATION REPORT NUMBER(S)			
6a. NAME OF PERFORMING ORGANIZATION University of Illinois Department of Computer Science	6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION Office of Naval Research, Code 1142PT Personnel & Training Research Programs	
6c. ADDRESS (City, State, and ZIP Code) 1304 W. Springfield Urbana, Illinois 61801		7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy St. Arlington, Virginia 22217-5000	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION	8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-85-K-0559	
8c. ADDRESS (City, State, and ZIP Code)		10. SOURCE OF FUNDING NUMBERS	
		PROGRAM ELEMENT NO 61153N	PROJECT NO. PR04206
		TASK NO. RR042060A	WORK UNIT ACCESSION NO NR667-551
11. TITLE (Include Security Classification) Learning Physical Domains: Toward a Theoretical Framework Unclassified			
12. PERSONAL AUTHOR(S) Kenneth D. Forbus, Dedre Gentner			
13a. TYPE OF REPORT Technical Report	13b. TIME COVERED FROM 85-9-1 TO 88-8-31	14. DATE OF REPORT (Year, Month, Day) 86 December	15. PAGE COUNT 38
16. SUPPLEMENTARY NOTATION			
17. COSATI CODES		18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD 05	GROUP 08	cognitive modelling; analogical reasoning qualitative reasoning; artificial intelligence	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) This paper presents a theoretical framework that is being developed in an attempt to construct a computational account of human learning of physical domains. Qualitative Process theory is used to model portions of people's physical knowledge, and structure-mapping theory is used to characterize the computations that move a learner from one representation to another. The paper outlines the component theories and proposes a learning sequence for physical domains.			
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT <input checked="" type="checkbox"/> UNCLASSIFIED//UNLIMITED <input type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS		21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Susan Chipman		22b. TELEPHONE (Include Area Code) 202-696-4218	22c. OFFICE SYMBOL ONR 1142PT

LEARNING PHYSICAL DOMAINS: TOWARD A THEORETICAL FRAMEWORK

Kenneth D. Forbus
Qualitative Reasoning Group
Department of Computer Science

Dedre Gentner
Psychology Department

December 1986

Department of Computer Science
University of Illinois at Urbana-Champaign
1304 W. Springfield Avenue
Urbana, Illinois 61801

This paper is a revised version of a paper that appears in
R. Michalski, J. Carbonell & T. Mitchell (Eds.),
Machine learning: An Artificial Intelligence Approach II,
Morgan Kaufmann Publishers, 1986. It supersedes BBN Technical
Report No. 5699 of the same title, published May, 1984.

This research is supported by the Office of Naval Research, Personnel
and Training Research Programs, Contract No. N00014-85-K-0559.

Approved for public release; distribution unlimited.



Accession For	
NTIS CRA&I	<input checked="checked" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

1. Introduction	2
1.1. Overview	2
2. Qualitative Process Theory	3
3. Comparisons and structure-mapping	8
4. Structure-Mapping and Learning	15
5. Stages of Understanding	18
5.1. Stage 1: Protohistories	19
5.2. Stage 2: The Causal Corpus	20
5.3. Stage 3: Naive Physics	22
5.4. Stage 4: Expert Models	26
6. Summary	30
7. Acknowledgments	32
8. References	32

1. Introduction

People use and extend their knowledge of the physical world constantly. Understanding how this fluency is achieved would be an important milestone in understanding human learning and intelligence, as well as a useful guide for constructing machines that learn. Our purpose is to construct a computational account of human experiential learning in physical domains.

We are still at the stage of refining the questions rather than providing detailed answers. In many cases, there is no direct evidence for our claims. In other instances, support for the theory is obtained by combining evidence from several different areas, including developmental psychology, studies of learning, and other psychological research. No one of these is adequate by itself. When extrapolating from adult learning research, we must keep in mind that cases of pure experiential learning are rare in adult life; some sort of instruction or prior expectation is typically involved. Developmental research provides a good source of data, since much of young children's learning is truly from direct experience. Yet when developmental results are applied it must be remembered that children are not only learning, but also maturing. Therefore, in order to isolate and study experiential learning, the existing empirical findings must be examined, filtered, and carefully fitted together. Although space does not permit detailing all the relevant lines of evidence, we will try to give the reader some justification for our claims whenever possible.

The past few years has seen significant progress in machine learning. However, to construct programs that learn as well as (or better than) people do, it is important to understand how human learning works. Ultimately both psychological studies and direct computational experiments (i.e., constructing programs) will be necessary to provide a full account. To this end, we will try when possible to indicate how techniques developed in machine learning might be used to implement such programs.

1.1. Overview

A brief prolog may help to organize the material. Three key ideas underlie the theory: (1) the centrality of *physical processes* in mental models of science; (2) the importance of analogy in learning; and (3) the primacy of rich, contextually specific representations. The idea that the notion of process is central to human knowledge about physical domains is the chief tenet of Qualitative Process (QP) theory (Forbus 1981; Forbus, 1984). This is not to say that notions of process are there from the beginning. Rather, we hypothesize that a person's experiential knowledge of a domain begins as a collection of scenarios that describe particular phenomena, out of which is developed a vocabulary of processes that provide a notion of mechanism for the domain. The second key idea concerns the role of comparisons among related knowledge structures. We conjecture that much of experiential learning proceeds through spontaneous comparisons - which may be implicit or explicit - between a current scenario and prior similar or analogous scenarios that the learner has stored in memory. Structure-mapping theory (see Gentner 1980; Gentner, 1983) describes these kinds of comparisons.

The third idea is a rather paradoxical claim: in human processing, more is often easier.¹ Rich, perceptually based representations are acquired earlier in learning than sparse abstract representations. That is, early domain representations differ from more advanced representations of the same domain in containing more information, especially perceptual

¹It should be noted that psychologists by no means generally agree with this claim. Consequently, we will try to be fairly explicit in presenting evidence for this position.

information specific to the initial context of use and acquisition. A second aspect of the "more is easier" claim concerns comparisons: we suggest that, for humans, similarity comparisons are easier when there is more overlap between the two knowledge structures being compared.

On the basis of these three ideas, we propose a canonical learning sequence. The claim is that human experiential learning of physical domains can be viewed as a sequence of different mental models: (1) *protohistories*, (2) the *causal corpus*, (3) *naive physics*, and (4) *expert models*. Briefly, protohistories are rich, contextually specific, highly perceptual representations of phenomena, capturing expectations about typical phenomenological patterns – for example, "If I turn the key, the car will start." With the causal corpus, the expectation of mechanism enters; here the representation consists of simple statements that some sort of causal connection exists between variables – "If the car has no gas, it will not start." In the naive physics stage, processes are introduced to provide the mechanism underlying the causal corpus – "Gas must flow from the tank to the carburetor and mix with air so that the mixture can be ignited by the spark." The disparate local connections of the causal corpus are replaced with qualitative models organized around the notion of process. Finally, in the expert models stage, quantitative representations are created – for example, models of the effects of different mixtures of oxygen and gasoline.

In this paper we discuss our conjectures about these models and how a learner constructs one type of model from another. First, however, the component theories that underlie this framework are briefly summarized: Qualitative Process theory, which provides concepts needed to represent the models (particularly in the naive physics stage); and structure-mapping theory, which characterizes the kinds of computations that move the learner from one representation to another. Then the overall role of structure-mapping comparisons is examined in the progression from rich to sparse representations. With these foundations in place, the four stages of learning for physical domains we postulate are then described.

2. Qualitative Process Theory

The first requirement is a language in which to describe people's common sense knowledge about physical situations. People know about a great many kinds of physical changes: things move, collide, bend, break, heat up, cool down, flow and boil. Intuitively we think of these as *processes*. Qualitative Process theory attempts to formalize this notion of process to provide a common form for qualitative theories of dynamics. As will be clear later on, we do not believe that the first models people construct of a domain take the form of processes, nor even that people become knowledgeable enough to construct these models for every domain they experience. Nevertheless, some of the concepts of QP theory will be useful for describing models in other stages as well.

In QP theory, a physical situation is modelled as a collection of objects and relationships among them, with processes responsible for causing changes. The continuous parameters of an object, such as temperature and pressure, are represented by *quantities*. A quantity has two parts, an *amount* and a *derivative*. Amounts and derivatives are both *numbers*. The model to keep in mind for numbers is that of the reals, but it is important to note that in QP theory particular numerical values are never used. Instead, the value of a number is described in terms of its *quantity space* – a collection of inequalities that hold between it and other quantities. Figure 1 illustrates a quantity space for the level of liquid in a container. The quantity space is a useful qualitative representation because processes typically start and stop when inequalities between parameters change.

Figure 1 - A quantity space

A quantity space describes the value of a number by the inequality relationships that hold between it and other numbers. An arrow indicates that the number at its head is greater than the number at its tail. Thus $LEVEL(wa)$ is less than $LEVEL(wb)$ and greater than $BOTTOM(a)$, while $LEVEL(wb)$ and $TOP(a)$ are unordered.




Figure 2 illustrates a typical process, called LIQUID-FLOW. A process has five parts: *individuals*, *preconditions*, *quantity conditions*, *relations*, and *influences*. Roughly speaking, the individuals part describes where instances of a process might occur, the preconditions and quantity conditions tell when it will be acting, and the relations and influences describe what holds as a consequence of it acting. In more detail, for any collection of objects that matches the individual specifications there is a *process instance* which represents the potential for that process to occur between those individuals in a particular way. For example, there will be two instances of LIQUID-FLOW between the liquid in the containers of figure 1, each corresponding to flow in a particular direction.

A process instance is *active* whenever both its preconditions and its quantity conditions are true. The distinction between preconditions and quantity conditions is that quantity conditions can be determined within QP theory but preconditions cannot. Quantity conditions concern what inequalities hold and what other processes (or individual views, which are introduced below) are active. Preconditions concern any relevant factors other than quantity conditions, such as spatial boundaries. For example, in "traditional" physics we can solve equations to figure out how fast a ball will be moving when it hits the floor, but the equations will not tell us a priori where the floor is. Or, returning to the present example, if we know that all the valves in the fluid path between the two containers are open (i.e., the fluid path is aligned) then fluid will flow, but we cannot predict within QP theory when or if someone will walk by and turn off a valve. Because these factors still affect dynamical conclusions, preconditions must be explicitly represented.

Figure 2 - A typical process

This process specification describes a simple kind of liquid flow. It can occur between two contained liquids that are connected by a fluid path, whenever the path is aligned - that is, all valves in the path are open - and the pressure in the one taken as source is greater than the pressure in the contained liquid taken as destination. The quantity type AMOUNT-OF represents how much "stuff" there is in an object. The function A maps a quantity into the number which is its amount, a number, as opposed to AMOUNT-OF, which is a function that maps a piece of stuff into a quantity.

Process LIQUID-FLOW

Individuals:

source, a CONTAINED-LIQUID
dest, a CONTAINED-LIQUID
path, a FLUID-PATH, FLUID-CONNECTION(source, dest, path)

Preconditions:

ALIGNED(path)

Quantity Conditions:

$A[\text{PRESSURE}(\text{source})] > A[\text{PRESSURE}(\text{dest})]$

Relations:

Let flow-rate, diff be quantities
 $\text{diff} = \text{PRESSURE}(\text{source}) - \text{PRESSURE}(\text{dest})$
 $\text{flow-rate} \propto \text{diff}$

Influences:

$I+(\text{AMOUNT-OF}(\text{dest}), A[\text{flow-rate}])$
 $I-(\text{AMOUNT-OF}(\text{source}), A[\text{flow-rate}])$

Whenever a process instance is active, its influences and relations hold. The influences component of a process specifies its direct effects; the relations component describes other facts that are true while the process is active. The direct effects-called *direct influences*-take the form

$I+(Q, n)$ or $I-(Q, n)$

depending on whether n is a positive or negative contribution to the derivative of Q . If a quantity is directly influenced, its derivative will be the sum of all the direct influences on it. Returning to the description of LIQUID-FLOW, for example, we see that when an instance of LIQUID-FLOW is active there will be a positive influence on the amount of liquid in the destination and an equal, negative influence on the amount of liquid in the source.

The relations field can describe new individuals that are created by virtue of the process being active (such as the steam produced by boiling water) as well as properties needed by

representations outside of QP theory (such as the appearance of boiling water). An especially important kind of fact expressed in the relations component is functional dependency between quantities. Functional dependencies between quantities are expressed by

$$Q1 \propto_{Q+} Q2$$

(read "Q1 is *qualitatively proportional* to Q2," or informally, "Q1 q-prop Q2"), meaning there exists a function which determines Q1 and is strictly increasing in its dependence on Q2. \propto_{Q-} indicates that the dependence is strictly decreasing. Note that qualitative proportionalities express partial information, since the exact nature of the function relating the parameters is not known and the function may or may not depend upon other quantities.² If a quantity Q1 is functionally dependent on a quantity Q2, and Q2 is influenced by a process P, then we will say that P *indirectly influences* Q1; that is, when P is acting it can cause Q1 to change. If, for instance, the PRESSURE and LEVEL of a liquid are qualitatively proportional to the AMOUNT-OF the liquid, then LIQUID-FLOW will indirectly influence both PRESSURE and LEVEL because it directly influences AMOUNT-OF. It is important to note that the only way a quantity can change is if it is directly or indirectly influenced. This means one can reason by exclusion: If nothing is influencing the amount of fluid in a container, then it isn't changing, but if the amount is changing, something must be influencing it. No changes happen by themselves. Furthermore, we can trace the possible paths of influences in a situation and determine whether or not particular kinds of changes can occur.

Two other important types of descriptions should also be mentioned here. *Individual views* are descriptions used to represent both objects whose existence are subject to dynamical constraints and states of objects. "The water in a cup," for example, is described as a CONTAINED-LIQUID, (see figure 3) because we can get rid of it by reducing its amount to zero (perhaps by making it the source of an instance of LIQUID-FLOW). Another example is a model of a spring. Springs have three states-relaxed, compressed, or stretched-each of which can be modeled by individual views. Individual views are specified in the same way that processes are, in that they have individuals, preconditions, quantity conditions, and relations. However, they do not have direct influences; directly influencing quantities is the sole prerogative of processes.

The other kind of description is the *encapsulated history*. How an object changes through time is represented by its *history* (Hayes 1979b). Histories are annotated pieces of space-time; thus they are object centered, have finite spatial extent, and extend over time.³ As its name suggests, an encapsulated history is a schematized description of some fragments of histories for a collection of objects. Encapsulated Histories are useful as summaries of behavior and to directly describe phenomena that have not been accounted for by process descriptions. An example of the latter usage is describing collisions between moving objects. A very simple way to model collisions is to say that the very next thing that happens after, say, an object hits a wall is that its velocity reverses and it starts moving the other way. Given how rapidly collisions occur, this model is quite adequate for most purposes, and encapsulated histories allow us to write it this way.

² QP theory also provides ways to specify dependence on properties that are not quantities (such as shape, in relating the level of a liquid in a container to its volume) and to make stronger statements about functional relationships, such as "Q1 depends on Q2 directly, with no intervening parameters" and "Q depends on Q1 and Q2 and nothing else" when required for framing stronger hypotheses about a domain. However, precise specifications of functions (e.g., $\sqrt{Q1 + Q2}$) are not permitted.

³ By contrast, the classic situational calculus (McCarthy & Hayes, 1969) description of change consists of *situations* that describe the whole universe at some particular instant of time.

Figure 3. - A typical individual view

This typical individual view describes a piece of liquid in a container, using the ontology for liquids described in (Hayes 1979a). There is just "syntactic sugar" for stating that whenever the preconditions and quantity conditions are true, *g* will exist.

INDIVIDUAL-VIEW CONTAINED-LIQUID

Individuals:

c a CONTAINER
s a SUBSTANCE

Preconditions:

CAN-CONTAIN-SUBSTANCE(*c*, *s*)

QuantityConditions:

A[AMOUNT-OF-IN(*s*, *c*)] > ZERO

Relations:

THERE IS *g*, a PIECE-OF-STUFF
HAS-QUANTITY(*g*, AMOUNT-OF)
AMOUNT-OF(*g*) = AMOUNT-OF-IN(*s*, *c*)
HAS-QUANTITY(*g*, LEVEL)
LEVEL(*g*) \propto_{Q+} AMOUNT-OF(*g*)
HAS-QUANTITY(*g*, PRESSURE)
PRESSURE(*g*) \propto_{Q+} LEVEL(*g*)

A reasoner's theory of dynamics for a particular domain is characterized in terms of (1) a *process vocabulary* that describes the kinds of processes the reasoner believes can occur and (2) a *view vocabulary* that describes dynamical objects and relevant states of objects. All changes are assumed to be directly or indirectly caused by processes—the *sole mechanism* assumption—which provides a strong constraint on the form of dynamical theories. Importantly, the content of dynamical theories is not tightly constrained—incorrect theories can be expressed as easily (and sometimes more easily!) than correct theories. For example, versions of Newtonian, Aristotelian, and Impetus theories of motion have all been encoded using QP theory.

QP theory sanctions several basic deductions. For example, the kinds of processes that might occur in a situation can be determined by using the process and view vocabularies to construct instances representing the different possibilities. The collection of processes acting at any time characterizes "what is happening" then in that situation, and these processes can be found by evaluating the preconditions and quantity conditions for these instances.

Consider again the example in Figure 1. There will be two instances of the LIQUID-FLOW process, and since the level in *wb* is greater than *wa*, the LIQUID-FLOW instance representing flow from *wb* to *wa* will be active. By taking into account all of the influences on each quantity (called

resolving its influences), we can often determine the sign of its derivative. The sign of the derivative is important because it represents how the amount of the quantity is changing—increasing, decreasing, or constant. In this example there is only one process instance acting, which makes things simple. AMOUNT-OF(wb) is directly influenced, and since this influence is negative it will decrease. By the \propto_{Q+} statements in the CONTAINED-LIQUID description, LEVEL(wb) and PRESSURE(wb) will be indirectly influenced and thus will also decrease. Similarly, AMOUNT-OF(wa), LEVEL(wa), and PRESSURE(wa) will increase.

From the ways the quantities are changing we can determine how the process and view structures themselves might change, since they depend in part on the inequalities stated as quantity conditions. This computation is called *limit analysis*. In the example, two things might happen—the pressures in wb and wa might equalize, or AMOUNT-OF(wb) could become zero, thus ending wb's existence (the geometry of this example rules out the latter).

The basic deductions of QP theory can be combined to perform more complex reasoning tasks. Two examples of more complex deductions are qualitative simulation (Forbus, 1984) and measurement interpretation (Forbus, 1983; 1986). Qualitative simulation consists of performing limit analysis repeatedly. It is useful for making predictions: for instance, that boiling water in a sealed container could cause an explosion. Measurement interpretation provides a link between physical theories and observations; for example, it might be hypothesized that the level of fluid in a container is dropping because the fluid is flowing out somewhere. Measurements may be interpreted by searching through the space of process and view structures, looking for situations where the results of influence resolution match the observations and which can be woven together to form a temporally consistent pattern of behavior.

3. Comparisons and structure-mapping

So far we have considered how portions of a person's knowledge about the physical world might be represented. Let us now turn to the question of how such domain models might be learned. We conjecture that a major process in experiential learning is comparing the current situation with stored descriptions. Consider for example a person who has just moved to a cold climate and is learning to operate a furnace. Suppose that at first he wrongly believes that the house will get warm faster if the thermostat is set to a temperature higher than the desired temperature. (Kempton (1985) shows that this view is quite common.) How can he reach the correct conclusion that the *rate* of heating does not depend on the temperature setting? There are at least three different ways, each based on a different kind of implicit comparison. First, he could compare his past furnace experiences with each other and notice a regularity in the rate of heating that is independent of the thermostat setting. Second, he may compare the furnace situation with known abstractions, and realize that it is best described as a position action controller (as opposed to a proportional-action controller). Third, he may use an analogy, comparing the furnace situation with a description from another domain, such as fluid flow, to suggest governing principles. Each of these ways of learning relies on some form of comparison, either with a stored record of literally similar events, with a stored abstraction, or with a stored description that can function as an analogy.

Structure mapping theory is concerned with such comparisons (see Gentner 1980, 1982, 1983; Gentner & Gentner, 1983). The theory describes the rules that are used to import a descriptive structure from one domain (the *base* domain) into another (the *target* domain). The central intuition is that an analogy suggests that a predicate structure from one domain can be applied in another domain with arbitrarily different objects and surface appearances. *Literal similarity, analogy, mere appearance mappings* and *abstraction mappings* (applications of general

laws) are viewed as different kinds of mappings between descriptions. The types of comparisons are defined syntactically, in terms of the form of the knowledge representation, not in terms of its content. Each type of comparison will be considered in turn.

1. An *analogy* is a comparison in which relational predicates, but few or no object attributes, are mapped from base to target. The particular relations mapped are determined by *systematicity*, as defined by the existence of higher-order constraining relations which can themselves be mapped.⁴ Thus, a relational chain – such as a causal chain – in the base that matches a relational chain in the target constitutes good support for its members. Winston (1983) gives an insightful demonstration of the need for such importance-dominated matching. The correspondences between objects of the base and objects of the target are determined by the roles of the objects in the relational structure, not by any intrinsic similarity between the objects themselves.

2. A *literal similarity* statement is a comparison in which a large number of predicates, both attributes and relations, can be mapped from base to target. Here, the model is based on one proposed by Tversky (1977), which states that the similarity between A and B increases with the size of the intersection of their feature sets and decreases with the size of the intersection of the two complement sets.⁵ There are many more shared predicates than nonshared predicates.

3. An *abstraction mapping* is a comparison in which the base domain is an abstract relational structure. Predicates from the abstract base domain are mapped into the target domain. As in analogy, the mapped predicates are a relational structure. Abstraction differs from analogy in the nature of the base domain. There are almost no object attributes in the base, so there are few, if any, one-place predicates to be left behind. Applying a rule to a situation is an example of abstraction mapping. Sometimes the relational structure so mapped will also be referred to as an *abstraction*.

4. A *mere-appearance* match is a comparison in which the object attributes match, but the relational structure does not. In a sense it is the opposite of analogy. Such matches are easily made; but they guarantee nothing beyond similarity in appearance.

A series of related examples, starting with the analogy between heat flow and water flow, will illustrate these distinctions. Figures 4a and 4b show a water-flow situation and the corresponding heat-flow situation (adapted from Buckley, 1979, pp. 15–25). Figure 5 shows a possible representation a person might have of the water situation. Notice that the description contains both object-attribute predicates, such as CYLINDRICAL(beaker), and relational predicates, such as


⁴ Object attributes are predicates which take one object as an argument, such as RED(x). We define the *order* of a proposition as follows. Constants and objects have order zero. The order of a proposition is one plus the maximum of the orders of its arguments. Thus CYLINDRICAL(x, y) would be first order if x and y are domain objects, and CYLINDRICAL(CYLINDRICAL(x, y), CYLINDRICAL(x, y)) would be second-order. Examples of higher-order relations are DANCE and LOVE.

⁵ Again, according to Tversky, the negative effects of the two complement sets are not equal; for example, given the question "How similar is A to B?" the set (B – A) – features of B not shared by A – counts more than the set (A – B).


Figure 4 Two Physical Situations Involving Flow

We will use these physical situations to illustrate the kinds of comparisons sanctioned by structure-mapping theory, and later to illustrate how QP-style domain descriptions can be used in analogies. Part (a) is a water-flow situation; part (b) is the corresponding heat-flow situation.

(a)



(b)



GREATER-THAN[PRESSURE(water, beaker), PRESSURE(water, vial)].

Let us consider the comparison types as exemplified here:

1. The analogy *Heat is like water* conveys that certain aspects of the water description can be mapped onto the heat domain. In particular, (1) object attributes should be dropped; (2) some relational predicates should be carried over; and (3) systematicity determines which relations should be mapped. Thus,

CYLINDRICAL (beaker)

is dropped, along with other object attributes; that is, the target objects do not have to resemble their corresponding base objects. Some relations are carried across, such as,

GREATER-THAN [PRESSURE (water, beaker), PRESSURE (water, vial)].

Yet not all relations are carried across. By the systematicity principle, this GREATER-THAN[PRESSURE(water, beaker), PRESSURE(water, vial)] relation is preserved because it is part of the mappable chain governed by the higher-order relation IMPLIES. In contrast, the relation

GREATER-THAN[CROSS-SECTIONAL-AREA(beaker), CROSS-SECTIONAL-AREA(vial)] is not carried across, since it is not part of any mappable system of constraining relations in this representation of the base domain.

Figure 6 shows the representation in the target domain of heat-flow that results from the analogical mapping. Given the object correspondences heat/water, beaker/coffee, vial/ice cube, pipe/bar, and PRESSURE/TEMPERATURE,⁶ systematicity operates to enforce a tacit preference for coherence and predictive power. The systematic relational structure in the water domain

IMPLIES(GREATER-THAN[PRESSURE(water, beaker),
PRESSURE(water, vial)],
FLOW(water, pipe, beaker, vial))

is mapped into

IMPLIES(GREATER-THAN[TEMPERATURE(heat, coffee),
TEMPERATURE(heat, ice cube)],
FLOW(heat, bar, coffee, ice cube)).

Note that the systematicity principle requires a *mappable* relational chain. If a particular chain of higher-order relations in of the base chain is not valid in the target, then another chain is selected. For example, suppose that we keep the same base domain – the system of containers shown in Figure 5 – but change the target domain. Suppose the two target objects are identical in temperature, but differ in their specific heats: say, a metal ball-bearing and a marble of equal mass. Now, the natural analogy concerns capacity differences in the base, rather than pressure differences. This is because the deepest relational chain that can be mapped to the target now concerns the situation in which pressures are equal in the base domain (analogously to temperatures being equal in the target domain):

IMPLIES[GREATER-THAN[CROSS-SECTIONAL-AREA (beaker),
CROSS-SECTIONAL-AREA (vial)],
GREATER-THAN [AMOUNT-OF-WATER (beaker),
AMOUNT-OF-WATER (vial)]]

This carries over into the target as

IMPLIES[GREATER-THAN[HEAT-CAPACITY (marble),
HEAT-CAPACITY (ball-bearing)],
GREATER-THAN[AMOUNT-OF-HEAT (marble),
AMOUNT-OF-HEAT (ball-bearing)]].

That is, given the same height (pressure) the container with a larger area will hold more water. Analogously, at the same temperature the object with greater heat capacity will hold more heat. Thus the interpretation of an analogy depends on the best relational match between base and target.

⁶ In this analogy, the first-order predicate of PRESSURE in the water domain must be replaced by TEMPERATURE in the heat domain. Although systems of relations can often be imported into the target without change, substitutions of functions, as well as objects and their attributes, are sometimes made in order to permit mapping a larger systematic chain.

Figure 5 - A representation of the water situation

This network represents a portion of what a person might know about the water situation illustrated in the figure 4. In this and other figures, predicates are written in upper case and circled. Objects are written in lower case and uncircled. A simplified representation is used to illustrate the rules of analogy. A more detailed model will be shown later.

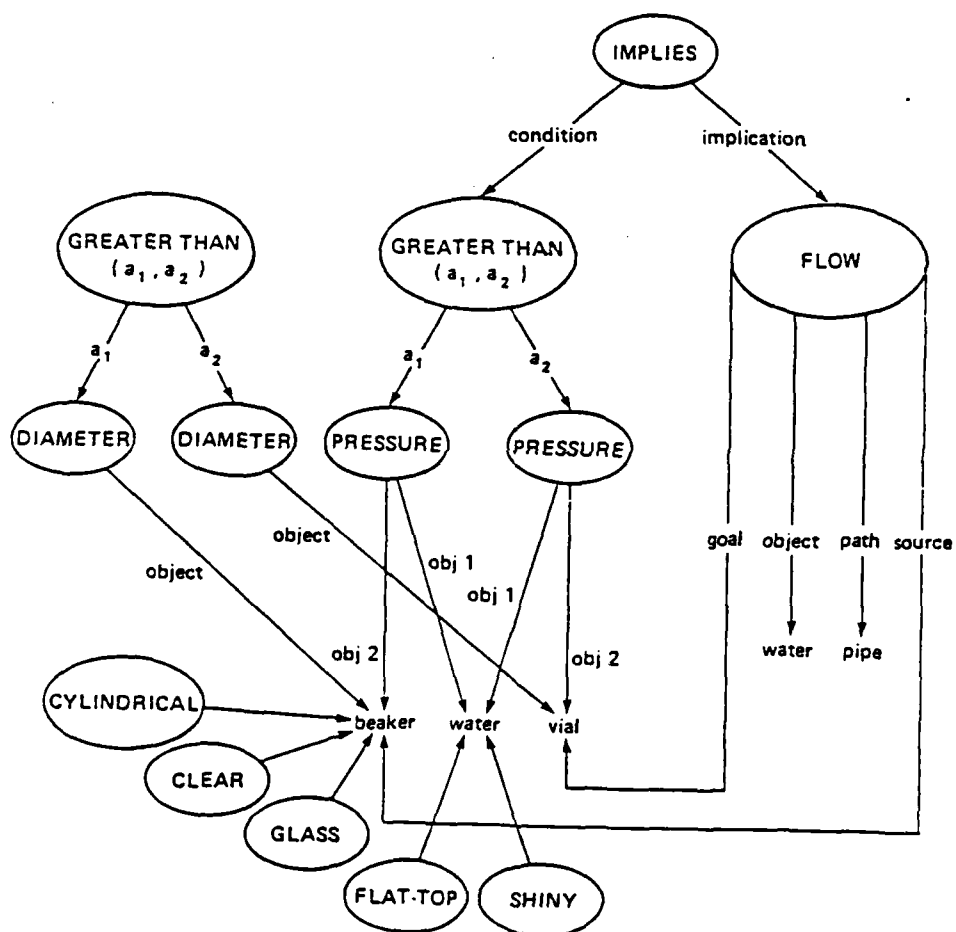
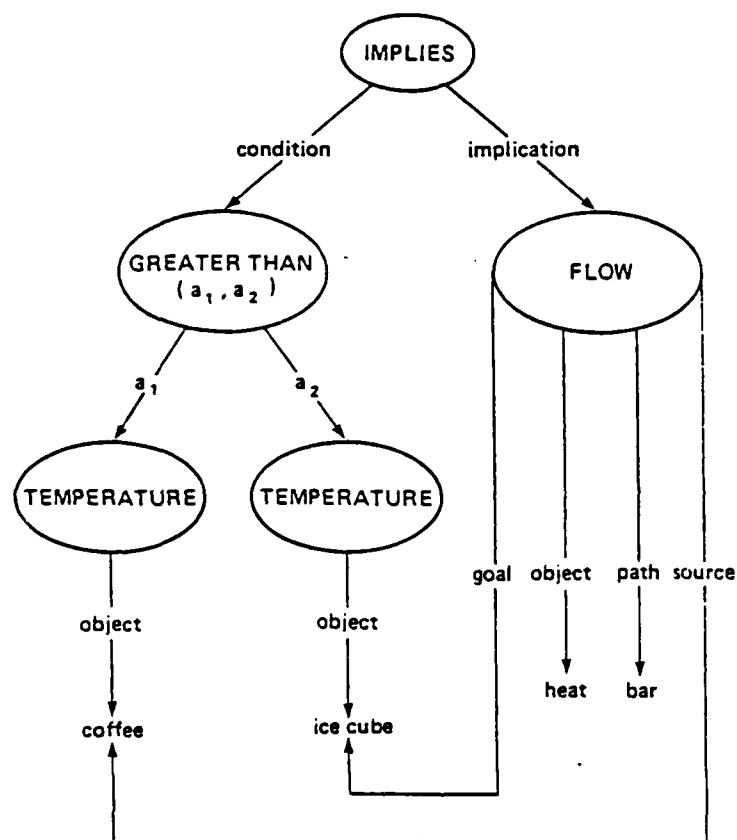


Figure 6 - A representation of the heat situation that results from the heat/water analogy. This network represents the knowledge a person would map across into the heat domain from the water situation illustrated in the previous figure. As in that figure, a simplified representation is used here. A more detailed treatment of this analogy is presented later.



2. The literal similarity comparison *Kool-Aid is like water* conveys that most of the water description can be applied to Kool-Aid. In literal similarity, both object attributes, such as *WET(water)*, and relational predicates, such as the systematic chain discussed above, are mapped over.

3. The abstraction *Heat is a through-variable* might be available to a student who knows some system dynamics. This abstraction conveys the idea that heat can be thought of as something that flows across a difference in potential (i.e., across some sort of "across-variable" - in this case, temperature). This is much the same relational structure as conveyed by the

analogy above; the difference is that in the abstract base domain of through-variables and across-variables there are no concrete properties of objects to be left behind in the mapping.

4. A mere-appearance match is a match with overlap chiefly in lower-order predicates, such as object-attributes, but little or no relational match. An example is *The tabletop sparkled like water*. Such a match typically yields little or no useful information about the target; here, for example, little can be learned about the table by mapping across knowledge about water. These matches, however, cannot be ignored in a theory of learning because a novice learner may be unable to tell them from true literal similarity matches.

Table 1 summarizes the kinds of predicate overlap that characterize literal similarity, analogy, abstraction, and mere appearance matches, as well as one other kind of comparison, *anomaly*. An anomaly is a match with little or no predicate overlap; it is included simply for completeness.

It should be clear that the contrasts described here are continua, not dichotomies. For example, analogy and literal similarity lie on a continuum. Given that two domains overlap in relational structure, then the comparison becomes more a literal similarity match to the extent that their object attributes also overlap, and more an analogy to the extent that few or no object attributes overlap. A different sort of continuum exists between analogies and general laws. In both cases, a relational structure is mapped from base to target. If the base representation includes concrete objects whose individual attributes must be left behind in the mapping, the comparison is an analogy. As the object nodes of the base domain become more abstract and variable-like, the comparison is seen as a general law.

4. Structure-Mapping and Learning

The role of a comparison in learning depends on at least two things: (1) *accessibility* - the likelihood that the match will be noticed, and (2) *usefulness* - what can be deduced from the match if it is accessed. Accessibility, in turn, depends at least on (a) the *familiarity* of the base description and (b) the overall *similarity* between the base description and the current target. The immediate usefulness of a match depends, of course, on whether the content of the match is

Table 1 Types of Comparisons

	ATTRIBUTES	RELATIONS	EXAMPLE
Literal Similarity	Many	Many	Milk is like water
Analogy	Few	Many	Heat is like water
Abstraction	Few	Many	Heat is a through variable
Anomaly	Few	Few	Coffee is like the solar system
Mere Appearance	Many	Few	The glass tabletop gleamed like a pool of water

appropriate to the task at hand. In addition, the usefulness of a match depends on the *inspectability* of the matching content – the degree to which it can be consciously analyzed and articulated. The comparisons discussed above behave very differently with respect to accessibility and inspectability.

For novice learners, literal similarity matches are the most accessible comparisons and abstractions, because they are typically extremely unfamiliar, are the least accessible. In contrast, abstraction matches are far more inspectable than literal similarity matches. On both dimensions, analogies are intermediate. This is one reason that analogy is crucial in learning: it is the novice's best route to an abstract, inspectable data structure. Some evidence for these conjectures will now be reviewed.

Surface matches are highly accessible. This includes both literal similarity matches and mere-appearance matches. It has been shown in education and training literature that the more similar the new situation is to the original situation the more readily transfer of training occurs (cf. Brown & Campione, 1985; Ross, 1984). The term "generalization gradient" expresses the fact that a learned response generalizes more readily the more similar the new situation is to the original situation. In contrast, subjects are often quite slow to use an available analogy. In research done by Reed, Ernst & Banerji (1974) and later by Gick and Holyoak (1980, 1983), subjects were asked to solve a rather difficult problem, such as how to cure an inoperable tumor with radiation without killing the flesh along the path of the rays. Just prior to receiving the problem some of the subjects read material that contained an analogous solution, such as a story about a general who split his troops up so that they all converged simultaneously on a fortress he wished to capture. There are three interesting results here. First, a good analogy can be very powerful, if it is noticed. Without the analogy, only about 10% of the subjects could solve the problem. Once the experimental subjects were told to use the prior story as an analogy, 80 to 90 percent of them solved the problem correctly. Second, a potentially powerful analogy can easily go unnoticed. Before the analogy was pointed out, only about a third of the subjects spontaneously noticed and used it. It cannot be taken for granted that a potential analog will be spontaneously noticed and used. Third, literal similarity is far more accessible than true analogy. In one of their studies, Gick and Holyoak (1983) happened to set up a literal similarity match between the story and problem. Subjects had to solve a problem involving tying two ropes together, and the story they were given involved tying two ribbons together. In this case, 70 to 80 percent of the subjects were able to access the matching story spontaneously. In a systematic series of studies, Ross (1984) varied the surface similarity between problems subjects were taught and later problems they had to solve and found that subjects were much more likely to be reminded of problems with similar surface content.

There is developmental evidence that literal similarity and mere-appearance matches appear prior to analogies and abstraction matches in learning. Kemler (1983) has found that young children group objects on the basis of overall similarity in situations where adults would group more analytically, using a single dimension. Keil and Batterman (1984) compared children's meanings and adults' and found a "characteristic-to-defining shift." For example, in defining "island" preschoolers use such surface features as "having palm trees and beaches" or "a warm place." Adults rely on defining features such as "surrounded by water." Another example occurs in labeling early word learning. In spontaneous labeling, one-year old children frequently apply words to objects that closely resemble the original referent of the word: for example *doggie* will be applied to another dog or to a cat, and *car* to cars, trucks or other vehicles (Clark, 1973). Truly analogous usages are seldom heard until the age of two or three years, when for example, a three-year old child might remark about his dirty, bedraggled blanket, "It's out of gas" (Gentner

& Stuart, 1984, Winner 1979).

Children are said to move from rich, concrete representations to more abstract, rule based systems (Bruner, Olver & Greenfield 1966; Gibson, 1969). Even three-year-olds can sort objects into perceptually similar categories; for example, they can group a cat and a dog and exclude a hen. However, not until they are five or six years can they succeed with a more abstract category, such as "living thing" which requires grouping perceptually dissimilar things. In the same vein, research on the novice-expert shift in adult learning has demonstrated that whereas novice science students typically match situations on the basis of surface features, experts use deeper and more abstract criteria (Larkin, 1983). For example, Chi, Feltovich, and Glaser (1981) have shown that when novice physics students are asked to classify problems into similar groups they put together problems with similar surface features, such as "inclined planes" or "pulleys." Experts, on the other hand, use categories like "force problems" and "energy problems."

One final indication of the ease with which literal similarity matches are made involves an indirect, but very important, line of argument. In the realm of object concepts, there is some evidence that people automatically perform literal similarity comparisons among perceptually similar experiences. Such comparisons are thought to result in composite prototypes (see Posner & Mitchell, 1967; Rosch, 1973, 1975, 1978; Smith & Medin, 1981).⁷ In the Posner & Mitchell study, people classified dot patterns into categories. After they had sorted the patterns, they were asked to remember which patterns they had seen. Although the task simply called for accessing verbatim memory, subjects showed systematic misrecognitions: they falsely remembered having seen prototypical patterns that were never presented. Thus, without being told to do so, people formed composite mental representations, apparently based on implicit comparisons among the patterns that they saw. Even theories which rely exclusively on stored exemplar information (such as that of Medin and Schaffer, 1978) share the assumption that literal similarity matches are made easily – indeed, automatically. The difference is that they assume that these implicit comparisons are made at the time of use of the stored exemplars, rather than assuming that the exemplars are encoded into a composite prototype. The virtually automatic nature of basic category learning is further evidence that the literal similarity matches on which they are based are highly accessible – indeed, evidence that making such comparisons is a passive, essentially automatic process (see also Reber 1967a, 1967b).

However, prototypes also illustrate the limited usefulness of literal similarity matches. Although these implicit composites are often sufficient for recognizing and categorizing situations, they are of limited use in deriving causal principles. This is because (1) a match based largely on perceptual commonalities will often fail to contain the correct principles, and (2) even when some of the correct relations are present, literal similarity matches are too rich to be inspectable. There is some evidence, albeit indirect, for this notion of rich, noninspectable representations. Nickerson and Adams (1979) studied people's memory of the common penny. Despite the overwhelming amount of experience that the subjects have had with pennies, and despite their evident ability to recognize and categorize pennies, they were remarkably poor at recalling or recognizing, given close near-misses, the details of how pennies look. This demonstrates that possessing a description sufficient to pick out a class of objects in ordinary life is no guarantee that the description can be articulated, or that it is very precise.

Studies of young children show that overall similarity judgments can be difficult to decompose. As discussed above, young children appear to base their similarity judgments on

⁷ The term *prototype* has been used in various ways in psychology. Here it is used to refer to a *structured* composite.

some kind of overall similarity (Kemler, 1983). Indeed Shepp (1978) has found that three- and four-year-olds are typically unable to judge one dimension independently of another. For example, they cannot ignore height when judging width. Unlike adults, they are unable to treat length and width as separable.

Abstraction matches are at the opposite pole from literal similarity. An abstraction match is likely to be extremely useful, in both respects: it should contain the correct principle, and the match should be inspectable. But abstractions are often not particularly accessible, especially for novices. Novice learners may not know the appropriate abstraction, or it may be so unfamiliar that they will not retrieve it when appropriate. Thus abstraction mappings, while ultimately important, are unlikely to play a major role in the early stages of learning.

Analogies lie between the highly accessible literal similarity matches and the highly useful abstraction matches. Potential analogies are less accessible in experiential learning than literal similarity matches (Gentner & Landers, 1985; Ross, 1984). This is because analogy requires accessing the learner's data base via relational matches; object matches are of little or no use. However, once found, an analogy should be more useful than a literal similarity match in deriving the key principles, since the shared data structure is sparse enough to permit analysis. (Of course, educators often explicitly introduce analogies in teaching beginners for exactly this reason. In this case, the problem of noticing the analogical match is bypassed. Moreover, by the systematicity principle, the set of overlapping predicates is likely to include higher-order relations, such as causality and logical implication. Thus analogy can function to reveal principles in a domain that previously lacked the appropriate abstractions (Burstein, in press; Carbonell, 1981, in press; Clements, 1982; Darden, in press; Gentner 1980, 1982; Gentner & Gentner, 1983; Gick & Holyoak, 1983; Hoffman, 1980; Van Lehn & Brown, 1980)

The *Analogical Shift Hypothesis* (Gentner, 1983) concerns the role of these comparisons in experiential learning. In the earliest stages most of the spontaneous matches are either mere-appearance matches (and thus erroneous) or are literal similarity matches, based on massive feature overlap. This is to say that initial learning is conservative, based on rich, specific-case kinds of matches. As the domain becomes familiar, more distant comparisons begin to occur; matches in which fewer object attributes are shared. These sparse comparisons lead to the kinds of binary connections that form the bulk of the causal corpus - for example, "lighter things go farther when thrown." Analogy also serves as a means of introducing structured mental models. Successful analogies may yield abstractions which can be stored and accessed (Gick & Holyoak, 1983). Winston's system (see Winston 1980; Winston 1982), which derives if-then rules by abstracting the predicates common to two analogs, shows how this can be done. Thus, analogy plays an important role in the middle and later stages of learning. In the final stages, when learning is well advanced, abstraction mappings play a major role.

5. Stages of Understanding

We suspect that four kinds of mental models are generated in the process of understanding physical domains. The sequence of models proposed here is developmental, in that the theories of each stage are generated both by the phenomena being understood and by the theories of the stage before it. It is not proposed that every person goes through each stage for every domain, nor that a person is at the same stage in every domain at the same time.

5.1. Stage 1: Protohistories

Suppose some new physical phenomenon is being observed. If there is no prior model, all one can do is observe and remember what is happening. We conjecture that the simplest physical models of a domain are *protohistories* – prototype histories which serve as summaries of experience.⁸ Like object prototypes, protohistories are the “most typical instances” of phenomena. The terms in these descriptions are observables, and their deductive import can be roughly expressed as “If I see X, then Y will happen (has happened).”

Consider a balance beam or seesaw. If a weight is placed on each side of the fulcrum, the seesaw will either tilt counterclockwise, tilt clockwise, or not tilt at all. Most people have had enough experiences with seesaws to have formed protohistories concerning their behavior. By the conjecture described here, a protohistory is automatically available whenever they encounter a seesaw. From it, they can often predict which way the particular seesaw will move. For example, they may have a protohistory that describes what happens if a small person gets on the seesaw opposite a large person.

However, the predictive power of protohistories is quite limited. There is no guarantee that the features matched actually correspond to relevant factors. For example, an observer will be fooled when a large person sits close to the fulcrum if the observer’s see-saw protohistories have been formed from watching people sitting at equal distances. Massive overlap in features is needed for reliable use, which means protohistories will yield conclusions in fewer situations than an abstract theory would. Consider, for example, two weights hung from opposite ends of a stick that is suspended by a string. The principle involved is the same, yet the situations look dissimilar enough that the protohistories for seesaws would not match. Furthermore, there is no certain way to decide between conflicting results if more than one protohistory matches a situation.⁹

5.1.1. Learning Protohistories

The process of constructing protohistories involves dividing up experience into classes according to literal similarity and abstracting a summary for each class. There has been little direct research on this process. However, investigations into the process of constructing object prototypes provides some hints. First, people seem to be able implicitly (i.e., unconsciously) to compute a kind of component match. Second, this intersection is not merely a simple feature intersection; rather, it appears that configurations among features are important in the prototype. Third, once this prototype is computed, it has powerful effects on subsequent processing of experience. As mentioned previously, once people abstract a prototype from a set of patterns they may be more confident of having seen the prototype – which was never presented – than they are of having seen the patterns actually presented (Posner 1967). Finally, people may not be aware of forming prototypes, except as a general sense of increased familiarity with a category.

To summarize, if protohistories behave like object prototypes, then they should be found to (1) be computed implicitly; (2) act as composite concepts; (3) be sensitive to perceptual configurations among events; and (4) once computed, show the recognition strength and other psychological privileges of prototypes.

⁸ Some of diSessa’s *phenomenological primitives* (diSessa 1983) appear to be representable as protohistories.

⁹ There are at least heuristic criteria, such as using the protohistory that has worked most often. The problem with such a criterion is that little is learned from mistakes.

The machine learning research that most closely captures this type of learning is concerned with *conceptual clustering* (see Michalski & Stepp 1983, 1983b). So far, such research has focused on classifying objects that can be characterized mainly by differing attributes. Extending such techniques to describe situations that depend critically on relational descriptions could provide a method for computing protohistories.

5.2. Stage 2: The Causal Corpus

Protohistories summarize the phenomena, but they do not constitute a theory of it. Building a detailed theory directly can be quite difficult. The space of possible models connecting all observable (and possible) parameters in a typical situation can be quite large. We conjecture that weaker theories, theories that characterize which parts of the situation are relevant to desired conclusions are formed first. In particular, we conjecture that a collection of CAUSE statements, the *causal corpus*, is computed from prototype objects and protohistories.

CAUSE is viewed here as an approximation concept, a weak form of ontological commitment. In particular, saying

CAUSE(A, B)

expresses belief in the existence of some mechanism, specified by some theory T, such that

$A \wedge T \rightarrow B$

Many, perhaps most, of the causal corpus relations are binary relations among variables – for example, “Bigger objects weigh more.” (Piaget, 1951; Smith, Wiser & Carey, in press), or “Smaller objects have higher pitch when struck.” (diSessa, 1983).

The notion of mechanism in the causal corpus is quite primitive: the causal beliefs need be neither explicit nor internally consistent. Later in the learning sequence, as we will see, *processes* will assume the role of mechanisms for physical domains. Nevertheless, we conjecture that, even at this early stage, the learner makes a distinction between mechanistic connections and, say, definitional connections.¹⁰ Further, we suspect that many of the initial causal connections are incorrect. Novices often include diagnostic and correlational relations in their causal corpus. For example, when asked if an increase in the evaporation rate will cause a change in the temperature of the water, a novice may reply “Yes, because it would have to be hotter to evaporate more.” But however vague or confused, a causal attribution is a statement of belief in some mechanistic connection.

The distillation of experience from protohistories into the causal corpus serves three purposes. First, it serves as a means of data reduction. Second, it provides a collection of heuristics that can be used directly to draw inferences. Even if the learner doesn’t have firm grounds to consider the CAUSE statements complete or correct, CAUSE statements may often suffice for the desired class of inferences. Third, the collection of CAUSE relations can be used to guide the search for a deeper theory of the domain. The CAUSE statements suggest connections among various aspects of the domain which a deeper theory must either explain or explain away.

Returning to the seesaw example, suppose the causal corpus is now applied to a balance beam built out of blocks. Suppose the two blocks on it are called a and b. The causal corpus might be as follows:

CAUSE(BIGGER(a, b), TILT-TOWARDS(a))

¹⁰ For example, the statement below is not a legitimate use of CAUSE by our account, since the required axioms of geometry do not specify a mechanism.

CAUSE(TRIANGLE(f), HAS-THREE-SIDES(f))

CAUSE(FARTHER(a, b), TILT-TOWARDS(a))

These statements can be interpreted as rules in several ways: If we see that block a is bigger than block b, one can predict tilt, and if one sees tilt, one may hypothesize that one block is farther out than another. These statements are more broadly applicable than protohistories since they refer to fewer properties. Unlike protohistories, the causal corpus is sparse enough to be debugged to some degree.

However, the approximate nature of the CAUSE relation and the binary characteristic of the laws limit the ability to discriminate between conflicting predictions. With the causal corpus above, for instance, if block a is bigger and block b is farther out, we will have two predictions. Inhelder & Piaget (1958) and Siegler (1976, 1981) have documented such a stage in the development of understanding about the balance beam (with analogous developmental sequences in other domains). Typically, children's first causal approach to the problem is to focus on weight. But there is an interesting second stage when they come to realize that both weight and distance are important but do not yet know the interrelations. They can manage either property by itself if the other is constant; but if both properties vary, they tend to focus on one or the other inconsistently. It is as though they had two separate binary laws. Eventually, they become able to coordinate weight and distance in the balance beam problem. At this stage, they have gone beyond the causal corpus. As will be discussed, in order to make more precise inferences the learner must eventually uncover the mechanisms whose behavior is described by causal corpus.

5.2.1. Learning the Causal Corpus

We suspect that there are several techniques for computing and debugging a causal corpus. The simplest technique is to hypothesize causality from co-occurrence, using rules like:

If you always see A before B, then hypothesize CAUSE(A, B)

and

If A is true whenever B is true, then hypothesize CAUSE(A, B)

These rules make certain assumptions on the form of memory, namely that some number of circumstances can be remembered, and that they can be remembered in sufficient detail that A and B are either explicitly stored or computable from what is stored. Protohistories should serve as a means of initial data reduction from which a causal corpus can be constructed.

It is not clear exactly how the learner abstracts out particular variables from the rich representations of the protohistory stage. One interesting mechanism is suggested by Medin and Wattenmaker's (in press) extension of the context theory (Medin & Schaffer, 1978). They suggest an abstraction mechanism whereby a similarity match which leads to correct predictions results in *common* information being augmented; whereas if a similarity match gives wrong predictions, the *differences* are augmented. However this is done, the simplification achieved with the causal corpus is considerable. Another study by Siegler (1978) shows the power of focusing on particular variables. Three-year-old children were shown a balance beam, asked to predict which way it would tilt, and then shown what actually occurred. Even after large numbers of trials, their performance failed to improve. But when they were taught to think of the domain in terms of a few relevant variables - weight and length - their performance did improve with experience. The moral to be drawn is that the pace of learning is greatly accelerated when a small number of variables can be abstracted from all the possibly relevant factors.

As suggested earlier, many of the early causal relations will be incorrect. We suspect that there exists a class of rules which are used to debug a causal corpus in the face of new information (c.f. Sussman, 1976). Each rule corresponds to a hypothesis about a bug in the structure of the causal corpus, such as a missing precondition. We believe that the task of

judging a causal corpus for consistency is an example of an important, but relatively neglected, kind of learning, *coherence-driven learning*. Coherence-driven learning is learning that is driven not by a mismatch between the model and the world but by inconsistencies within the model itself. Williams, Hollan, & Stevens (1983) found evidence of such learning. They studied a subject who was learning about a heat exchanger, and noted that one source of insight was a "boggle" experience, in which the person noticed that a current inference contradicted a prior belief. We are still examining the criteria for judging the consistency of a causal corpus.¹¹ Such criteria will play a major role in controlling the debugging rules and in the mixture of generation and debugging that occurs.

Analogy provides another important technique for extending a causal corpus (see Gentner & Gentner, 1983; Stevens, 1979). The CAUSE relations from one domain can be mapped into another, since CAUSE qualifies as a higher-order constraining relation (see also Winston, 1982).

5.3. Stage 3: Naive Physics

The naive physics models replace CAUSE statements with theories about the specific mechanisms of change. The ontology is extended by adding processes to explain observed changes. The ontology also includes properties and objects that are not directly observable (for example, heat and heat flow) and the new relationships (such as fluid path and heat path) required to reason about them.

An important advantage of these models is the ability to reason by exclusion. In the naive physics stage, unlike the previous stages, predictions that fail still yield information about the situation. For instance, if fluid is flowing into a container and the level is not rising, then it is reasonable to hypothesize that fluid is flowing out of it through some unknown path.

Returning to our balance beam, a process SWING might be used to describe rotation around a contact point (See figure 7). The preconditions describe the geometric configuration of the system, and the quantity condition says that SWING will occur whenever there is a non zero angular velocity. SWING directly influences the angular position of the beam. Thus a prediction concerning tilt becomes a prediction about which instance, if any, of the SWING process will be active.¹²

What influences ANGULAR-VELOCITY? The existence of an ANGULAR-ACCELERATION process (see Figure 8) that directly influences ANGULAR-VELOCITY whenever there is a net torque will be assumed. It is further assumed that

$$\begin{aligned} & \forall(x) \forall(y) \\ & \text{PHYSOB}(x) \wedge \text{CONTACT-POINT}(cp) \\ & \rightarrow \text{NET-TORQUE}(x, cp) = \text{SUM-OF}(\text{TORQUES-ON}(x, cp)) \end{aligned}$$

In other words, the net torque on an object around a contact point is the sum of the torques on that object measured about that contact point. The mass of the beam will be ignored, and pull

¹¹ With Lance Rips, of the University of Chicago, we are investigating the role of intransitives in debugging causal representations as one form of coherence-driven learning.

¹² An alternate, and equally good, representation for ω with ω would leave directions implicit in the sign of the angular velocity. In that vocabulary, the balance beam would give rise to only one instance of ω with ω and determining whether ω is the beam moving requires that one determine first whether the instance of swing is active and if it is, what the sign of the angular velocity is.

Figure 7 - A SWING process

A SWING process describes rotation of an object around another object. For the balance beam there will be two instances of this process, differing only in their bindings for the direction dir. In each instance b will be bound to the beam, c will be bound to the fulcrum, and cp will be bound to the contact point between them.

It is assumed that each physical object (PHYSOB) has quantities to represent its angular position and velocity with respect to each point of contact with other objects. Directions will be noted by the symbols CW, CCW, and NULL, corresponding to clockwise rotation, counterclockwise rotation, and no rotation.

Process SWING

Individuals:

b a PHYSOB
c a PHYSOB
cp a CONTACT-POINT
dir a DIRECTION

Preconditions:

MOBILE(b)
not MOBILE(c)
CONNECTED(b, c, cp)
ROTATION-FREE(b, c, cp)
DIRECTION-OF(dir, ANGULAR-VELOCITY(b, cp))

Quantity Conditions:

Am[ANGULAR-VELOCITY(b, cp)] > ZERO

Influences:

I+(ANGULAR-POSITION(b, cp), A[ANGULAR-VELOCITY(b, cp)])

Figure 8 - ANGULAR-ACCELERATION process
Process ANGULAR-ACCELERATION

Individuals:

b a PHYSOB
c a PHYSOB
cp a CONTACT-POINT
dir a DIRECTION

Preconditions:

MOBILE(b)
not MOBILE(c)
CONNECTED(b, c, cp)
ROTATION-FREE(b, c, cp)
DIRECTION-OF(dir, NET-TORQUE(b, cp))

Quantity Conditions:

Am[NET-TORQUE(b, cp)] > ZERO

Relations:

Let acc be a quantity
 $acc \propto_{Q+} NET-TORQUE(b, cp)$
 $acc \propto_{Q-} MASS(b)$

Influences:

I+(ANGULAR-VELOCITY(b, cp), A[acc])

of gravity on the blocks on each side of the fulcrum will be assumed to be the only source of torques. Figure 9 describes this induced torque by means of an individual view. Notice that the factors illuminated in the causal corpus of BIGGER and FARTHER have become the quantities MASS and DISTANCE, and their role in the producing swinging has been explicated. In particular, these properties determine how much torque each block places on the beam. The sum of the torques determines the net torque, which can cause the beam to accelerate and thus swing.

This model comes one step closer to a model that can always determine which way something will tilt. There will still be cases in which exactly what will happen cannot be determined (e.g., if the mass on one side is increased and it is brought closer to the pivot), but this is a precise hypothesis about what all the relevant factors are.

5.3.1. Learning Naive Physics

The major problem in learning a naive physics is constructing a vocabulary of processes that consistently describes experience. The learner must strip away the irrelevant predicates that are part of the protohistories and causal corpus and construct more appropriate

Figure 9 - A description of gravity-induced torque

Positive torques are assigned to clockwise (CW) and negative torques are assigned to counter-clockwise (CCW).

Individual View GRAVITY-INDUCED-TORQUE

Individuals:

b a PHYSOB
c a PHYSOB
d a PHYSOB
cp a CONTACT-POINT

Preconditions:

CONNECTED(b, c, cp)
ON(d, b)

Relations:

Let f be a quantity
 $f \in \text{TORQUES-ON}(b, cp)$
 $f \propto_{q+} \text{DISTANCE}(C-M(d), cp)$
 $f \propto_{q+} \text{MASS}(d)$
 $\text{ON}(C-M(d), \text{SIDE-OF}(CW, b, cp)) \text{ iff } As[f] = 1$
 $\text{ON}(C-M(d), \text{SIDE-OF}(CCW, b, cp)) \text{ iff } As[f] = -1$
 $\text{ON}(C-M(d), \text{SIDE-OF}(NULL, b, cp)) \text{ iff } As[f] = 0$

descriptions. In addition, the learner must sometimes hypothesize the existence of objects and properties that are not directly observable. Research in machine learning, particularly the work on *explanation-based learning* (Mitchell, 1986; DeJong & Mooney, 1986), should be useful here. Several researchers are already beginning to directly address such problems in modelling scientific discovery (Langley 1983; Falkenhainer, 1985; Rajamoney, DeJong, & Faltings, 1985).

The causal corpus provides a search space for potential process vocabularies. Each statement in the causal corpus must be elaborated into a consequence of a process vocabulary. It appears that there is only a small number of distinct ways to perform the elaboration, depending on the particular form of the arguments. For example, the statement

The decrease in AMOUNT-OF q causes the LEVEL of q to fall.

indicates that some active process (or individual view) in the situation contains the statement

$\text{LEVEL}(q) \propto_{q-} \text{AMOUNT-OF}(q)$

in its relations.

Hypothesizing a process vocabulary from a causal corpus should be much simpler than working from protohistories or direct observation. Yet it still appears difficult. We conjecture that there are several constraints that make the problem more tractable. First, people are apparently *conservative* in the introduction of unobservable properties. For example, some subjects have a model of a domain that appears to be organized around one parameter – a “generalized strength” attribute. In reasoning about fluids, for instance, they appear to use *pressure*, *flow rate*, and *velocity* as different names for the same thing. In electricity, they use *voltage*, *current*, *power*, *potential*, and *velocity of electrons* interchangeably. The advantage of this generation strategy is, of course, that simpler models will be explored first, with further distinctions made only when necessary. Second, some physical laws are used as constraints on what process vocabularies are possible. Conservation of energy, for example, demands that if a process directly influences a quantity representing some form of energy, it must also directly influence some other quantity representing some form of energy, but in the opposite direction.

Once again, analogy can provide a constructive mechanism. It can be used to import candidate processes from previously understood domains – for example, as when one understands electricity in terms of water flow (Gentner & Gentner, 1983) or evaporation in terms of an implicit model of rocket ships escaping from earth (Collins & Gentner, in press). This is an especially powerful mechanism because if the model for the previous domain is consistent with physical laws, then it suggests that the model for the new domain may be so as well.

We can illustrate this with an analogy from liquid-flow to heat-flow. Recall the liquid flow model presented in Section 2. Figure 10 illustrates a collection of assertions which describe the consequences of a particular instance of LIQUID-FLOW.¹³ Suppose a person hypothesizes that there is a process of heat flow which is analogous to the process of liquid flow. By the structure-mapping theory, this means that the person suspects that a similar relational structure holds among the objects in the heat-flow situation (the coffee, the ice cube, the silver bar, and the instance of heat flow) as among the objects in the liquid-flow situation (the water in the beaker, the water in the vial, the pipe, and the instance of liquid flow). Mapping the systematic relational structure (see Figure 11), leads to several predictions that the person can check to see whether the analogy is correct. For example, it can be determined whether or not the temperature of the ice cube is rising and the temperature of the coffee falling. The structure-mapping rules for analogy have provided an initial model for the process of heat flow; in particular, the preconditions, quantity conditions, relations, and influences are all carried across from liquid flow. Note that to make the analogy really work, a new kind of object – a HEAT-PATH – must be postulated. Thus analogy can provide candidates for extending ontologies.¹⁴

5.4. Stage 4: Expert Models

The models generated so far have two important limitations. First, they still contain fundamental ambiguities, ambiguities which are inherent in the nature of qualitative

¹³ The assertions were generated by an early version of GIZMO, a computer program constructed to explore the computational aspects of QP theory. GIZMO was designed to make predictions and interpret measurements, not to be a learning system. In particular, these descriptions were not generated with learning or analogy in mind.

¹⁴ Of course, such extensions are not to be made lightly. The authors suspect that new types of objects are postulated in the target domain only when necessary to preserve a much larger systematic structure.

Figure 10 – Relational structure for an instance of liquid flow

Depicted below are several important conclusions which follow from the definition of liquid flow presented in Figure 2 and the assumption that an instance of liquid flow exists involving the liquids in the two containers. Specifically, they describe the conditions for and consequences of the process instance p1-0 being active.

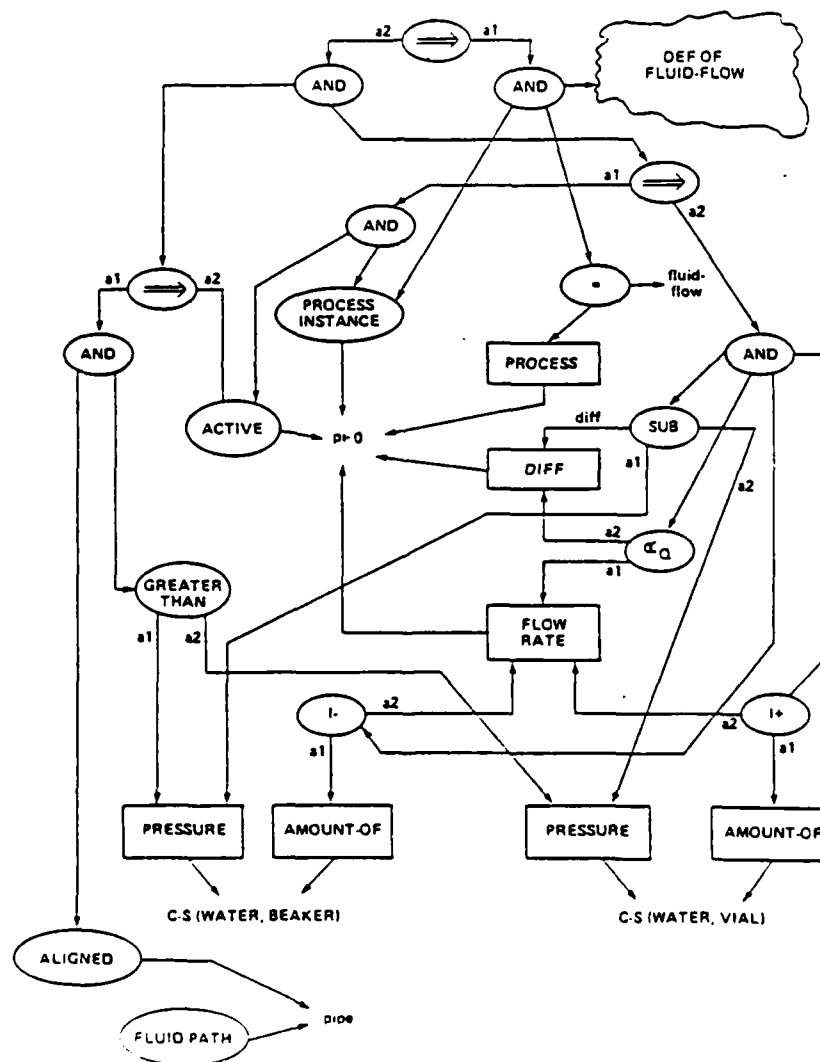
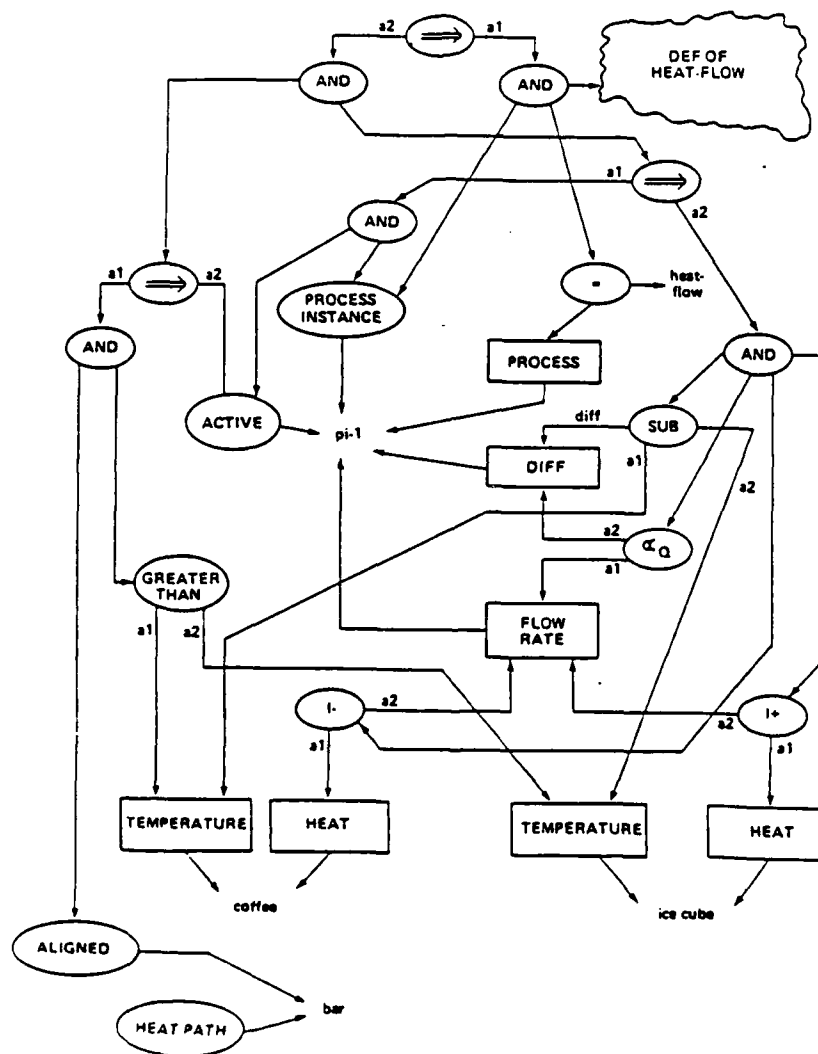


Figure 11 – Relational structure transferred to heat flow

Here the relational structure describing a situation involving liquid flow has been transferred to a situation involving heat flow. Notice the systematicity of the relational structure, as indicated by the nested chains of implications.



representations.¹⁵ Second, they lack domain-independent generalizations (except in the raw form of the representation - CAUSE statements, processes, and so on). The final stage of learning consists of overcoming these limitations, of discovering ways to resolve ambiguities and to construct powerful generalizations.

Clearly several kinds of knowledge are involved, and the potential complexity of the models in this stage is open-ended (it includes the whole of mathematical physics, for example). Examples of the kinds of knowledge involved include equations to describe relationships between parameters, "rules of thumb" to specify useful default resolutions for ambiguities, and new ontologies to allow reasoning about more complex systems. The importance of mathematical models is fairly obvious. The rules of thumb are less obvious but equally important (see e.g., Lenat, 1982). In physical domains they include empirical knowledge about the circumstances under which certain processes can be ignored (such as evaporation when water is poured from one glass to another) and what their net effect is (such as Black's law for the temperature of mixtures). Finally, different ontologies are sometimes necessary to deal with certain types of complex systems. In the process-oriented physics discussed here, describing flow requires finding flow paths. Finding flow paths in complex networks such as electrical circuits can quickly become computationally intractable; switching to a device-centered physics (such as that described in deKleer & Brown, 1983) can reduce the computational burden to manageable proportions for such systems.

To complete the balance beam example, we know that the force of a block on the beam is qualitatively proportional to the mass of the block and to the distance from the fulcrum. If we also know that the torque is the product of distance and weight, then providing numerical values for these quantities will allow an unambiguous prediction about tilt.

5.4.1. Learning Expert Models

The transition to expert models involves several kinds of learning. Some aspects of this transition probably lie outside the scope of experiential learning; for example, people typically learn mathematical models by being taught rather than by discovery. Some aspects of this learning - such as developing new ontologies - involve improving the content of the representations. Other aspects of the transition from a naive physics to an expert physics are better described as translating the existing qualitative representations into quantitative statements, using mathematics to express laws. By converting a physical theory into a mathematical model, the learner gains the ability to make precise predictions and to recognize powerful generalizations more easily. An important part of this refinement is to elaborate Q^+ statements into constraint equations. Langley (1979; Langley, Zytkow, Simon & Bradshaw, 1983) and Falkenhainer (1985) describe techniques that should be useful for converting qualitative laws into mathematical relations.

Developing rules of thumb means knowing not just what is possible, but what is probable. The learner must discover which outcomes raised by qualitative reasoning are likely or unlikely and which potential interactions can be ignored. The techniques developed in machine learning for acquiring heuristics should be directly applicable (c.f. Lenat, 1982; Mitchell, 1981). In addition, the authors suspect the possible behaviors raised by naive physics are compared against known protohistories. Hypothesized outcomes that have no corresponding protohistory are judged unlikely, and those corresponding to a highly familiar and accessible protohistory are

¹⁵ The nature of ambiguity in qualitative descriptions is discussed by de Kleer & Brown (1983) and Forbus (1984).

judged very likely (see Tversky & Kahneman, 1973).

Further, it seems likely that at least some expert rules of thumb derive from learning new and better protohistories. This intuition is based in part on research in automaticity (Schneider & Fisk, 1983). It has been demonstrated that, given an orderly domain and sufficient practice, adult subjects can learn a new response pattern well enough so that it becomes essentially effortless (see also Anderson, 1982; Rumelhart & Norman, 1978). Moreover, there is some transfer from this over-learned material to new similar material. These learned sequences have many of the essential qualities of protohistories. First, they are triggered by recognition (in the terms used here, by a literal similarity match between the present situation and a stored situation). Second, computing and carrying out the procedures that follow from the match is automatic; virtually no attentional resources are required. Third, these computations are implicit; subjects are typically poor at introspecting about what they are doing, and when they do introspect, it can interfere with the response (Brooks, 1978; Reber, 1967, 1976). It may be too simplistic to view protohistories as a special case of automatic pattern-response combination. Nevertheless, there is enough overlap in the phenomena to allow some confidence that protohistories can continue to be learned at all stages of expertise. Of course, the *contents* of expert protohistories may be different from those of novices, since experts' protohistories may reflect a more advanced ontology, as discussed below. However, the mechanism of a perceptually-triggered automatic match should be the same.

We suspect that ontological shift is driven both by the desire to understand more complex physical systems and by the emergence of domain-independent mathematical abstractions. As an example of the first kind, consider the problem of reasoning about fluid flow in a complex system, such as a steam plant. Hayes (1979b) has distinguished two separate ontologies for liquids: a *contained-liquid* ontology, in which liquid is thought of as the fluid in a place and a *molecular collection* ontology, in which water is thought of as little bits of fluid that move around inside the system. The contained-liquid ontology is appropriate if the goal is to determine what flows can occur. However, it will not help us determine how changes in the properties of the working fluid in one part of the system (say, the rising temperature of the inlet water in a boiler) can affect properties of the fluid in another part of the system (say, temperature of the steam coming out of the boiler's superheater). In this case, liquid must be viewed in terms of molecular collections that move around inside the system. Conversely, establishing flows using the molecular collection view is very difficult. A learner with only one of these two ontologies will have a difficult time with certain questions, and such difficulties may drive the search for a new ontology.

Mathematical abstractions provide another important driving force in ontological change. In system dynamics, for example, physical systems involving fluid elements, mechanical elements, thermal elements, and acoustical elements are viewed as variations on a common, abstract theme. This means that the analysis and synthesis tools developed for abstract mathematical models can be used to solve problems in several domains. This is a powerful motivation, as evidenced by the wave of interest in attempting diverse applications evoked by the publications of certain new mathematical formalisms (e.g., catastrophe theory and fractal geometry).

6. Summary

We have described our progress in combining structure mapping theory and Qualitative Process theory into a framework that aims to account for experiential learning in physical

domains. The learning sequence is built around three ideas. First, development proceeds from rich to sparse, and from concrete to abstract—that is, initial representations differ from later representations in containing more information, and in particular, more context-specific information. Second, after sufficient experience, people develop experiential models that are centered around the notion of physical process, as described by Qualitative Process theory. Third, implicit processes of comparing and mapping between stored knowledge and a current situation, as described in structure-mapping theory, are central to experiential learning.

Four stages of experiential learning have been laid out: protohistories, the causal corpus, naive physics and expert models.¹⁶ The first stage, that of protohistories, embodies the idea that early representations are rich and context specific; this stage attempts to capture a combination of evidence from developmental patterns, similarity judgments, basic-level categories and object prototypes. The third stage is the process-centered stage described by Qualitative Process theory. The fourth stage builds on the third stage models, adding domain-independent generalizations and in some cases mathematical models. There is some evidence for the third and possibly the fourth stages in the research on expertise under the rubric of the novice-expert shift (Chi, Feltovich and Glaser, 1981; Larkin, 1983).

The second stage, the causal corpus, is the most speculative. There is no direct evidence for its existence, nor do we currently have a detailed theory of the kinds of causal statements that can enter into the representations. Moreover, detailing how the causal corpus emerges from protohistories will not be easy. But something like the causal corpus seems necessary: a collection of simplistic, mostly binary directed regularities among dimensions and quantities that begin to be differentiated out of the tangled representations of the protohistory stage. The learner can use these simple assertions as grist for further progress.

What happens to prior stages as new stages occur? First, stored representations have to be distinguish from new learning. We conjecture that learners retain much of their stored knowledge even when they go beyond the stage at which it was formed. Thus, a hydraulics engineer still uses the same protohistory he or she formed as a child to decide how fast to carry a glass of water without spilling it. And, as de Kleer (1979) points out, expert physicists and engineers do not always resort to quantitative models (fourth stage); frequently the answer they want can be obtained by using a good qualitative model (third stage).

But what about new learning? Does new learning occur only at the leading edge, or do people continue to learn at levels below the most advanced stage they have attained? We suspect that even experts continue to learn at all prior stages, with the possible exception of the causal corpus. As described earlier, there is evidence that experts continue to lay down new protohistories. Similarly, learners who are operating at the fourth stage, that of expert models, will continue to learn and refine their naive physics. This is because the mathematical models of the fourth stage are not a substitute for the process models of the third stage.¹⁷ Improvement in domain physics are useful whether or not mathematical models are also available. As experts

* We are using the term *protohistory* for the first stage. We have been influenced by Piaget's well-known distinction between the *concrete* and *abstract* stages of development (see Piaget, 1963). However, we are using the term *protohistory* to refer to Piaget's *concrete* stage, and not to Piaget's *abstract* stage. This is because we believe that the *concrete* stage is the most important stage of development.

¹⁶ The fourth stage is the *expert models* stage. We have been influenced by Piaget's well-known distinction between the *concrete* and *abstract* stages of development (see Piaget, 1963). However, we are using the term *expert models* to refer to Piaget's *concrete* stage, and not to Piaget's *abstract* stage. This is because we believe that the *concrete* stage is the most important stage of development.

increases the least new learning is expected within the causal corpus.

Of the four levels, the causal corpus has the least claim to continued independent existence in an advanced expert. The causal corpus is not reliable for prediction, nor does it possess the advantages of automaticity.¹⁸ In summary, the overall picture is that a learner moves from rich, perceptually specific protohistories to the sparser representation of the causal corpus. The causal corpus serves as a staging area in which rough connections among variables can be stored until they can be subsumed into a true system. If learning continues, a person develops a process-centered naive physics, and, for some domains, expert models.

7. Acknowledgments

We gratefully acknowledge the intellectual and financial support of Schlumberger-Doll Research in our initial efforts. This research is supported by the Office of Naval Research, Contract No. N00014-85-K-0559. We thank Allan Collins, Dave Chapman, Ed Smith, Albert Stevens, and Dan Weld for their insightful comments on an earlier draft of this paper.

8. References

- Anderson, J. R., Acquisition of cognitive skill, *Psychological Review* 89(4):369-406, 1982
- Brooks, L., Nonanalytic concept formation and memory for instances, *Cognition and Categorization*, Lloyd, B. B. and Rosch, E., Eds., Erlbaum Associates, Hillsdale, NJ, 1978
- Brown, A. L., and Campione, J. C., Three faces of transfer: Implications for early competence, individual differences, and instruction, in M. Lamb, A. Brown, & B. Rogoff, *Advances in developmental psychology*, Erlbaum Associates, Hillsdale, N. J., in press
- Bruner, J. S., Olver, R., & Greenfield, P., *Studies in cognitive growth*, Wiley, New York, 1966
- Buckley, S., *Sun up to sun down*, McGraw-Hill Company, New York, 1979
- Burstein, M. H., Concept formation by incremental analogical reasoning and debugging. In *Proceedings of the International Machine Learning Workshop*, University of Illinois, Monticello, Illinois, June, 1983
- Carbonell, J. G., A computational model of problem-solving by analogy, *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, August, 1981
- Carbonell, J. G., Derivational analogy in problem solving and knowledge acquisition, In *Proceedings of the International Machine Learning Workshop*, University of Illinois, Monticello, Illinois, June, 1983

¹⁸ The discussion of course concerns domains for which the learner does eventually acquire expert knowledge. As a result, people rely heavily on causal corpus knowledge in domains in which they are inexpert. Further, there are domains, such as child rearing or getting rich, which lack definitive models. Collins' (1978) work on plausible reasoning suggests that in these domains, people rely heavily on the causal corpus knowledge. See also Salter's (1981) work on the causal corpus in domains.

Case, R., Intellectual development from birth to adulthood: A neo-piagetian interpretation. In Siegler, R. S., *Children's thinking: What develops?* Erlbaum Associates, Hillsdale, N.J., 1978

Chi, M. T. H., Feltovich, P. J., and Glaser, R., Categorization and representation of physics problems by experts and novices, *Cognitive Science*, 5(2), April-June, 1981

Clark, E. V., What's in a word? On the child's acquisition of semantics in his first language, In Moore, T. E., *Cognitive development and the acquisition of language*, Academic Press, New York, 1973

Clement, J., Analogical reasoning patterns in expert problem-solving, *Proceedings of the Fourth Annual Conference of the Cognitive Science Society*, Berkeley, 1982

Collins, A., Fragments of a theory of human plausible reasoning, In D. Waltz, *Theoretical issues in natural language processing II*, University of Illinois, Urbana, Illinois, 1978

Collins, A. & Gentner, D., How people construct mental models, In D. Holland & N. Quinn (Eds.) *Cultural models in language and thought*, Cambridge, England: Cambridge University Press, in press

Collins, A. Warnock, E.H., Aiello, N., & Miller, M.L., Reasoning from incomplete knowledge. In Bobrow, D. & Collins, A., eds, *Representation and Understanding: Studies in cognitive science*, Academic Press, New York, 1975

Darden, L., Reasoning by analogy in scientific theory construction, *Proceedings of the International Machine Learning Workshop*, University of Illinois, Monticello, Illinois, June, 1983

deJong, G. F. & Mooney, R. J., Explanation-Based Learning: An alternate view, *Machine Learning*, Volume 1, Number 2, April, 1986

de Kleer, J., *Causal and Teleological Reasoning in Circuit Recognition*, Technical Report No. 529, MIT Artificial Intelligence Laboratory, Cambridge, Massachusetts, September, 1979

de Kleer, J. and Brown, J.S., *Assumptions and Ambiguities in Mechanistic Mental Models*, In D. Gentner, & A. L. Stevens, (Eds.), *Mental Models*, Erlbaum Associates, Hillsdale, N.J., 1983

Dietterich, T. and Michalski, R. S., A comparative review of selected methods for learning from examples. In *Machine Learning*, Michalski, R. S., Carbonell, J., and Mitchell, T., eds., Tioga Publishing Company, Palo Alto, California, 1983

diSessa, A., Phenomenology and the Evolution of Intuition, In D. Gentner, & A. L. Stevens, (Eds.), *Mental Models*, Erlbaum Associates, Hillsdale, N.J., 1983

Falkenhainer, B. Proportionality graphs, units analysis, and domain constraints: Improving the power and efficiency of the scientific discovery process. *Proceedings of IJCAI 85*, Los Angeles, California, August 1985

Flavell, E. H., *The developmental psychology of Jean Piaget*, Van Nostrand Reinhold Company,

New York, 1963

Forbus, K., Qualitative reasoning about physical processes, *Proceedings of IJCAI-7*, Vancouver, B.C., August, 1981

Forbus, K., Measurement Interpretation in Qualitative Process Theory, *Proceedings of IJCAI-8*, 1983

Forbus, K., Qualitative Process Theory, *Artificial Intelligence* 24:85-169, 1984

Forbus, K., The problem of existence. In *Proceedings of the Seventh Annual Conference of the Cognitive Science Society*, Irvine, CA, August, 1985

Forbus, K. Interpreting observations of physical systems, technical report No. UIUCDCS-R-86-1248, Department of Computer Science, University of Illinois, March 1986. A short version will appear in the proceedings of AAAI-86.

Garner, W. R., Aspects of a stimulus: Features, dimensions, and configurations, In Lloyd, B.B. & Rosch, E., (Eds.), *Cognition and Categorization* Erlbaum Associates, Hillsdale, N.J., 1978

Gelman, R., Counting in the Preschooler: What does and does not develop, In Siegler, R. S., *Children's thinking: What develops?*, Erlbaum Associates, Hillsdale, N.J.,

Gentner, D., *The structure of analogical models in science*, Technical Report No. 4451, Bolt Beranek and Newman Inc., Cambridge, Mass, 1980

Gentner, D., Are scientific analogies metaphors?, In Miall, D., *Metaphor: Problems and perspectives*, Harvester Press, Ltd., Brighton, England, 1982

Gentner, D., Structure-mapping: A theoretical framework for analogy, *Cognitive Science* 7(2), 1983

Gentner D., & Gentner, D. R., Flowing waters or teeming crowds: Mental models of electricity, In D. Gentner & A. L. Stevens. (Eds.), *Mental Models*, Erlbaum Associates, Hillsdale, N.J., 1983

Gentner, D. & Landers, R., Analogical reminding: A good match is hard to find, *Proceedings of the International Conference on Systems, Man and Cybernetics*, Tucson, Arizona, November 1985

Gentner, D. & Stuart, P. Metaphor as structure-mapping: What develops? Technical Report No. 314, Center for the Study of Reading, University of Illinois, Urbana, IL, June 1984

Gibson, E. J., *Principles of perceptual learning and development*, Prentice-Hall, New York, 1969

Gick, M. L., & Holyoak, K. J., Analogical problem solving, *Cognitive Psychology*, 12:306-355, 1980

Gick, M. L., & Holyoak, K. J., Schema Induction and Analogical Transfer, *Cognitive Psychology*, 15(1):1-38, 1983

Hayes, P. J., The Naive Physics Manifesto, In Michie, D., *Expert systems in the micro-electronic age*, Edinburgh University Press, May, 1979

Hayes, P. J., *Naive Physics 1 - Ontology for Liquids*, Memo, Centre pour les etudes Semantiques et Cognitives, Geneva, Switzerland, 1979

Hesse, M. B., *Models and analogies in science*, University of Notre Dame Press, Notre Dame, Indiana, 1966

Hoffman, R. R., Metaphor in science, In Honeck, R. P. & Hoffman, R. R., *The psycholinguistics of figurative language*, Erlbaum Associates, Hillsdale, N.J., 1980

Inhelder, B., & Piaget, J., *The growth of logical thinking from childhood to adolescence*, Basic Books, New York, 1958

Kahneman, D. and Tversky, A., *Subjective probability: A judgement of representiveness jR*, *Cognitive Psychology*, 3:430-454, 1972

Keil, F. & Batterman, N. A., *A characteristic-to-defining shift in the development of word meaning*, *Journal of Verbal Learning and Verbal Behavior*, 1984, 23, p. 221-236.

Kemler, D. G. Holistic and analytic modes in perceptual and cognitive development. In T. J. Tighe & B. E. Shepp (Eds.), *Perception, cognition and development: Interactional analysis*. Hillsdale, New Jersey: Erlbaum, 1983, p. 77-102.

Kempton, W., Two theories used for home heat control., In D. Holland and N. Quinn (Eds.), *Cultural models in language and thought*., 1983

Langley, P., Rediscovering physics with BACON.3, In *Proceedings of IJCAI-6*, 1979

Langley, P., Zytkow, J., Simon, H. and Bradshaw, G., Mechanisms for Qualitative and Quantitative Discovery, *Proceedings of the Second International Machine Learning Workshop*, 1983

Larkin, J., The role of problem representation in physics, In D. Gentner & A. L. Stevens (Eds.), *Mental Models*., Erlbaum Associates, Hillsdale, N.J., 1983

Lenat, D., The Nature of Heuristics, *Artificial Intelligence*, 19:189-249, 1982

McCarthy, J. and Hayes, P., Some philosophical problems from the standpoint of artificial intelligence, *Machine Intelligence 4*, Edinburgh University Press, 1969

Medin, D. L. & Schaller, M. M., Context theory of classification learning, *Psychological Review*, 1978, 85, No. 3, 207-238.

Medin, D. L. & Wattenmaker, W. D. Category cohesiveness *Theories and cognitive archeology*, in press. Michalski, R. S., A Theory and Methodology of Inductive Learning, *Artificial Intelligence*, 20:111-161, 1983

Michalski, R. S., and Stepp, R. E., Learning from Observation: Conceptual Clustering. In Michalski, R. S., Carbonell, J. and Mitchell, T., *Machine Learning*, Tioga Publishing, Palo Alto, California, 1983

Michalski, R. S., and Stepp, R. E., How to Structure Structured Objects, *Proceedings of the International Machine Learning Workshop*, University of Illinois, Monticello, Illinois, 1983

Mitchell, T. M., Utgoff, P. E., Nudel, B. and Banerji, R., Learning problem-solving heuristics through practice, *Proceedings of the Seventh International Joint Conference on Artificial Intelligence*, August, 1981

Mitchell, T. M., Generalization as Search, *Artificial Intelligence*, 18:203-226, 1982

Mitchell, T. M., Keller, R., & Kedar-Cabelli, S., Explanation-Based Generalization: A unifying view, *Machine Learning*, Volume 1, Number 1, January, 1986

Nelson, K., How children represent knowledge of their world in and out of language: A preliminary report. In Siegler, R. S., *Children's thinking: What develops?*, Erlbaum Associates, Hillsdale, N.J., 1978

Nickerson, R. S., and Adams, M. J., Long-term memory for a common object, *Cognitive Psychology*, 11:287-307, 1979

Palmer, S. E., Fundamental aspects of cognitive representation, In Lloyd, B. B. & Rosch, E., (Eds.), *Cognition and Categorization*, Erlbaum Associates, Hillsdale, N.J., 1978

Piaget, J., *The child's conception of physical causality*, Routledge & Kegan Paul, London, 1951

Piaget, J., *The construction of reality in the child*, Basic Books, New York, 1954

Posner, M. I., & Mitchell, R. F., Chronometric analysis of classification, *Psychological Review*, 74:392-409, 1967

Rajamoney, S., DeJong, G. F., & Faltings, B., Towards a model of conceptual knowledge acquisition through directed experimentation, *Proceedings of IJCAI 1985*, Los Angeles, California.

Reber, A. S., Implicit learning of artificial grammars, *Journal of Verbal Learning and Verbal Behavior*, 6:855-863, 1967

Reber, A. S., Implicit learning of synthetic languages: The role of instructional set, *Journal of Experimental Psychology: Human Memory and Learning*, 2:88-94, 1976

Reed, S. K., Ernst, G. W., and Banerji, R., The role of analogy in transfer between similar problem states, *Cognitive Psychology*, 6:436-450, 1974

Rendell, L. R., Substantial constructive induction using layered information compression: Tractable feature formation in search (revised). *Department of Computer Science Report*

UIUCDCS-R 85-1198, University of Illinois, May 1985

Rendell, L. R., A general framework for induction and a study of selective induction. *Machine Learning*, in press.

Rosch, E., On the internal structure of perceptual and semantic categories, In Moore, T. E., *Cognitive development and the acquisition of language*, Academic Press, New York, 1973

Rosch, E., Cognitive representations of semantic categories, *Journal of Experimental Psychology: General*, 107:192-233, 1975

Rosch, E., Principles of categorization, In Rosch, E. & Lloyd, B. B., *Cognition and categorization*, Erlbaum Associates, Hillsdale, NJ, 1978

Ross, B., Reminders and their effects in learning a cognitive skill, *Cognitive Psychology*, 1984, 16, 371-416.

Rumelhart, D. & Norman, D. A., Accretion, tuning, and restructuring: Three modes of learning, In Cotton, J. W. & Klatzky, R. L., *Semantic Factors in Cognition*, Erlbaum Associates, Hillsdale, NJ, 1978

Schneider, W. & Fisk, A. D., Attention theory and mechanisms for skilled performance, In Magill, R. A., *Memory and control of action*, North-Holland Publishing Company, New York, 1983

Shepp, B. E., From perceived similarity to dimensional structure: A new hypothesis about perspective development, In Rosch, E. & Lloyd, B. B., *Cognition and categorization*, Erlbaum Associates, Hillsdale, N.J., 1978

Siegler, R. S., Three aspects of cognitive development, *Cognitive Psychology*, 8:481-520, 1976

Siegler, R. S., Developmental sequences within and between concepts, *Monographs of the Society for Research in Child Development*, 46(2), 1981

Siegler, R. S., The origins of scientific reasoning, In Siegler, R. S., *Children's thinking: What develops?*, Erlbaum Associates, Hillsdale, N.J., 1978

Smith, E. E., & Medin, D. L., *Categories and concepts*, Harvard University Press, Cambridge, Massachusetts, 1981

Stevens, A., Collins, A., & Goldin, S.E., Misconceptions in student's understanding, *Journal of Man Machine Studies*, 11:145-156, 1979

Sussman, G. J., *A computational model of skill acquisition*, Springer-Verlag, New York, 1976

Tversky, A., Features of similarity, *Psychological Review*, 84:327-352, 1977

Tversky, A., & Gati, I., Similarity, separability and the triangle inequality, *Psychological Review*,

89(2):123-154, 1982

Tversky, A., and Kahneman, D., Availability: A heuristic for judging frequency and probability, *Cognitive Psychology*, 5:207-232, 1973

Van Lehn, K. & Brown, J. S. Planning nets: A representation for formalizing analogies and semantic models of procedural skills. In R. E. Snow, P. A. Federico & W. E. Montague (Eds.) *Aptitude, learning and instruction: Cognitive process analyses*. Hillsdale, N. J.: Lawrence Erlbaum Associates, 1980.

Williams, M., Hollan, J., and Stevens, A. L., Human Reasoning about a simple physical system, In Gentner, D. and Stevens, A. L., (Eds.), *Mental Models*, Erlbaum Associates, Hillsdale, NJ, 1983

Winner, E., New names for old things: The emergence of metaphoric language. *Journal of Child Language*, 6:469-491, 1979

Winston, P. H., Learning and reasoning by analogy., *Communications of the ACM*, 23(12), 1980

Winston, P. H., Learning new principles from precedents and exercises, *Artificial Intelligence*, 19:321-350, 1982

Distribution List [Illinois/Gentner] NR 667-551

Dr. Phillip L. Ackerman
University of Minnesota
Department of Psychology
Minneapolis, MN 55455

Dr. Beth Adelson
Dept. of Computer Science
Tufts University
Medford, MA 02155

AFOSR,
Life Sciences Directorate
Bolling Air Force Base
Washington, DC 20332

Dr. Robert Ahlers
Code N711
Human Factors Laboratory
Naval Training Systems Center
Orlando, FL 32813

Dr. Ed Aiken
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. James Anderson
Brown University
Center for Neural Science
Providence, RI 02912

Dr. John R. Anderson
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Steve Andriole
George Mason University
School of Information
Technology & Engineering
4400 University Drive
Fairfax, VA 22030

Technical Director, ARI
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Gary Aston-Jones
Department of Biology
New York University
1009 Main Bldg
Washington Square
New York, NY 10003

Dr. Patricia Baggett
University of Colorado
Department of Psychology
Box 345
Boulder, CO 80309

Dr. Eva L. Baker
UCLA Center for the Study
of Evaluation
145 Moore Hall
University of California
Los Angeles, CA 90024

Dr. Meryl S. Baker
Navy Personnel R&D Center
San Diego, CA 92152-6800

prof. dott. Bruno G. Bara
Unità di ricerca di
intelligenza artificiale
Università di Milano
20122 Milano - via F. Sforza 23
ITALY

Dr. William M. Bart
University of Minnesota
Dept. of Educ. Psychology
330 Burton Hall
178 Pillsbury Dr., S.E.
Minneapolis, MN 55455

Leo Beltracchi
U.S. Nuclear Regulatory Comm.
Washington, D.C. 20555

Dr. Mark H. Bickhard
University of Texas
EDB 504 ED Psych
Austin, Texas 78712

Dr. Gautam Biswas
Department of Computer Science
University of South Carolina
Columbia, SC 29208

Dr. John Black
Teachers College, Columbia Univ
525 West 121st Street
New York, NY 10027

Dr. R. Darrell Bock
University of Chicago
NORC
6030 South Ellis
Chicago, IL 60637

Dr. Sue Bogner
Army Research Institute
ATTN: PERI-SF
5001 Eisenhower Avenue
Alexandria, VA 22333-5600

Dr. Jeff Bonar
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Gordon H. Bower
Department of Psychology
Stanford University
Stanford, CA 94306

Dr. Robert Breaux
Code N-095R
Naval Training Systems Center
Orlando, FL 32813

Dr. Ann Brown
Center for the Study of Reading
University of Illinois
51 Gerty Drive
Champaign, IL 61820

Dr. John S. Brown
XEROX Palo Alto Research
Center
3333 Coyote Road
Palo Alto, CA 94304

Dr. Bruce Buchanan
Computer Science Department
Stanford University
Stanford, CA 94305

Maj. Hugh Burns
AFHRL IDE
Lowry AFB, CO 80230-5000

Dr. Patricia A. Butler
OERI
555 New Jersey Ave., NW
Washington, DC 20208

Distribution List [Illinois/Gentner] NR 667-551

Dr. Joseph C. Campione
Center for the Study of Reading
University of Illinois
51 Gerty Drive
Champaign, IL 61820

Joanne Capper
Center for Research into Practice
1718 Connecticut Ave., N.W.
Washington, DC 20009

Dr. Jaime Carbonell
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Dr. Susan Carey
Harvard Graduate School of
Education
337 Gutman Library
Appian Way
Cambridge, MA 02138

Dr. Pat Carpenter
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

LCDR Robert Carter
Office of the Chief
of Naval Operations
OP-01B
Pentagon
Washington, DC 20350-2000

Chair, Department of
Psychology
College of Arts and Sciences
Catholic University of
America
Washington, DC 20064

Dr. Fred Chang
Navy Personnel R&D Center
Code 51
San Diego, CA 92152-6800

Dr. Davida Charney
English Department
Penn State University
University Park, PA 16802

Dr. Paul R. Chatelier
OUSDRE
Pentagon
Washington, DC 20350-2000

Dr. Michelene Chi
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Dr. L. J. Chmura
Computer Science and
Systems Branch
Naval Research Lab.
Washington, DC 20375-5000

Mr. Raymond E. Christal
AFHRL/MOE
Brooks AFB, TX 78235

Professor Chu Tien-Chen
Mathematics Department
National Taiwan University
Taipei, TAIWAN

Dr. Yee-Yeen Chu
Perceptronics, Inc.
21111 Erwin Street
Woodland Hills, CA 91367-3713

Dr. William Clancey
Stanford University
Knowledge Systems Laboratory
701 Welch Road, Bldg. C
Palo Alto, CA 94304

Dr. Charles Clifton
Tobin Hall
Department of Psychology
University of
Massachusetts
Amherst, MA 01003

Dr. Allan M. Collins
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Stanley Collier
Office of Naval Technology
Code 222
800 N. Quincy Street
Arlington, VA 22217-5000

Dr. William Crano
Department of Psychology
Texas A&M University
College Station, TX 77843

Brian Dallman
3400 TTW/TTGXS
Lowry AFB, CO 80230-5000

Dr. Laura Davis
NRL/NCARAI, Code 7510
4555 Overlook Ave., SW
Washington, DC 20375-5000

Dr. Natalie Dehn
Department of Computer and
Information Science
University of Oregon
Eugene, OR 97403

Dr. Gerald F. DeJong
Artificial Intelligence Group
Coordinated Science Laboratory
University of Illinois
Urbana, IL 61801

Goery Delacote
Directeur de L'informatique
Scientifique et Technique
CNRS
15, Quai Anatole France
75700 Paris FRANCE

Dr. Sharon Derry
Florida State University
Department of Psychology
Tallahassee, FL 32306

Dr. Andrea di Sessa
University of California
School of Education
Tolman Hall
Berkeley, CA 94720

Dr. R. K. Dismukes
Associate Director for Life Sciences
AFOSR
Bolling AFB
Washington, DC 20332

Dr. Stephanie Doan
Code 6021
Naval Air Development Center
Warminster, PA 18974-5000

Distribution List [Illinois/Gentner] NR 667-551

Dr. Emanuel Donchin
University of Illinois
Department of Psychology
Champaign, IL 61820

Dr. Jean Claude Falmagne
Department of Psychology
New York University
New York, NY 10003

Dr. John R. Frederiksen
Bolt Beranek & Newman
50 Moulton Street
Cambridge, MA 02138

Defense Technical
Information Center
Cameron Station, Bldg 5
Alexandria, VA 22314
Attn: TC
(12 Copies)

Dr. Beatrice J. Farr
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Norman Frederiksen
Educational Testing Service
Princeton, NJ 08541

Dr. Thomas M. Duffy
Communications Design Center
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. Pat Federico
Code 511
NPRDC
San Diego, CA 92152-6800

Dr. Michael Friendly
Psychology Department
York University
Toronto Ontario
CANADA M3J 1P3

Dr. Richard Duran
University of California
Santa Barbara, CA 93106

Dr. Paul Feltovich
Southern Illinois University
School of Medicine
Medical Education Department
P O Box 3926
Springfield, IL 62708

Julie A. Gadsden
Information Technology
Applications Division
Admiralty Research Establishment
Portsmouth, Portsmouth PO6 4AA
UNITED KINGDOM

Dr. John Ellis
Navy Personnel R&D Center
San Diego, CA 92252

Mr. Wallace Feurteig
Educational Technology
Bolt Beranek & Newman
10 Moulton St.
Cambridge, MA 02238

Dr. Michael Genesereth
Stanford University
Computer Science Department
Stanford, CA 94305

Dr. Susan Embretson
University of Kansas
Psychology Department
426 Fraser
Lawrence, KS 66045

Dr. Gerhard Fischer
University of Colorado
Department of Computer Science
Boulder, CO 80309

Dr. Dedre Gentner
University of Illinois
Department of Psychology
603 E. Daniel St.
Champaign, IL 61820

Dr. Randy Engle
Department of Psychology
University of South Carolina
Columbia, SC 29208

J. D. Fletcher
9931 Corsica Street
Vienna VA 22180

Chair, Department of
Psychology
George Mason University
Fairfax, VA 22030

Dr. Susan Epstein
Hunter College
144 S. Mountain Avenue
Montclair, NJ 07042

Dr. Linda Flower
Carnegie-Mellon University
Department of English
Pittsburgh, PA 15213

Chair, Department of
Psychology
Georgetown University
Washington, DC 20057

ERIC Facility-Acquisitions
4833 Rugby Avenue
Bethesda, MD 20014

Dr. Kenneth D. Forbus
University of Illinois
Department of Computer Science
1304 West Springfield Avenue
Urbana, IL 61801

Dr. Robert Glaser
Learning Research
& Development Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. K. Anders Ericsson
University of Colorado
Department of Psychology
Boulder, CO 80309

Dr. Barbara A. Fox
University of Colorado
Department of Linguistics
Boulder, CO 80309

Distribution List [Illinois/Gentner] NR 667-551

Dr. Arthur M. Glenberg
University of Wisconsin
W. J. Brogden Psychology Bldg.
1202 W. Johnson Street
Madison, WI 53706

Dr. Sam Glucksberg
Department of Psychology
Princeton University
Princeton, NJ 08540

Dr. Susan Goldman
University of California
Santa Barbara, CA 93106

Dr. Sherrie Gott
AFHRL/MODJ
Brooks AFB, TX 78235

Dr. T. Govindaraj
Georgia Institute of Technology
School of Industrial & Systems
Engineering
Atlanta, GA 30332

Dr. Wayne Gray
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. James G. Greeno
University of California
Berkeley, CA 94720

Dr. Dik Gregory
Behavioral Sciences Division
Admiralty Research Establishment
Teddington, Middlesex
ENGLAND

Dr. Gerhard Grossing
Atominstitut
Schuttenelstrasse 115
Vienna, AUSTRIA A-1020

Prof. Edward Haertel
School of Education
Stanford University
Stanford, CA 94305

Dr. Henry M. Half
Half Resources, Inc.
4918 33rd Road, North
Arlington, VA 22207

Dr. Ronald K. Hambleton
Prof. of Education & Psychology
University of Massachusetts
at Amherst
Hills House
Amherst, MA 01003

Stevan Harnad
Editor, The Behavioral and
Brain Sciences
20 Nassau Street, Suite 240
Princeton, NJ 08540

Dr. Wayne Harvey
SRI International
333 Ravenswood Ave.
Room B-5324
Menlo Park, CA 94025

Dr. Reid Hastie
Northwestern University
Department of Psychology
Evanston, IL 60201

Prof. John R. Hayes
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Barbara Hayes-Roth
Department of Computer Science
Stanford University
Stanford, CA 95305

Dr. Frederick Hayes-Roth
Teknowledge
525 University Ave.
Palo Alto, CA 94301

Dr. Shirley Brice Heath
School of Education
Stanford University
Stanford, CA 94305

Dr. Joan I. Heller
505 Haddon Road
Oakland, CA 94606

Dr. Jim Hollan
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. Melissa Holland
Army Research Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Keith Holyoak
University of Michigan
Human Performance Center
330 Packard Road
Ann Arbor, MI 48109

Ms. Julia S. Hough
Lawrence Erlbaum Associates
6012 Greene Street
Philadelphia, PA 19144

Dr. James Howard
Dept. of Psychology
Human Performance Laboratory
Catholic University of
America
Washington, DC 20064

Dr. Earl Hunt
Department of Psychology
University of Washington
Seattle, WA 98105

Dr. Ed Hutchins
Intelligent Systems Group
Institute for
Cognitive Science (C-015)
UCSD
La Jolla, CA 92093

Dr. Barbara Hutson
Virginia Tech
Graduate Center
2990 Telestar Ct.
Falls Church, VA 22042

Dr. Barbel Inhelder
University of Geneva
Geneva SWITZERLAND 12U-4

Distribution List [Illinois/Gentner] NR 007-551

Dr. Dillon Inouye
WICAT Education Institute
Provo, UT 84057

Dr. Alice Isen
Department of Psychology
University of Maryland
Catonsville, MD 21228

Dr. Robert Jannarone
Department of Psychology
University of South Carolina
Columbia, SC 29208

Dr. Claude Janvier
Directeur, CIRADE
Universite' du Quebec a Montreal
Montreal, Quebec H3C 3P8
CANADA

Dr. Robin Jeffries
Hewlett-Packard Laboratories
P O Box 10490
Palo Alto, CA 94303-0971

Dr. Robert Jernigan
Decision Resource Systems
5595 Vantage Point Road
Columbia, MD 21044

Margaret Jerome
c/o Dr. Peter Chandler
83, The Drive
Hove
Sussex
UNITED KINGDOM

Chair, Department of
Psychology
The Johns Hopkins University
Baltimore, MD 21218

Dr. Douglas A. Jones
Thatcher Jones Assoc.
P O Box 6640
10 Trafalgar Court
Lawrenceville
NJ 08648

Dr. Marcel Just
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Daniel Kahneman
The University of British Columbia
Department of Psychology
#154-2053 Main Mall
Vancouver, British Columbia
CANADA V6T 1Y7

Dr. Ruth Kanfer
University of Minnesota
Department of Psychology
Elliott Hall
75 E. River Road
Minneapolis, MN 55455

Dr. Mary Grace Kantowski
University of Florida
Mathematics Education
359 Norman Hall
Gainesville, FL 32611

Dr. Milton S. Katz
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Frank Keil
Department of Psychology
Cornell University
Ithaca, NY 14850

Dr. Wendy Kellogg
IBM T. J. Watson Research Ctr
P O. Box 218
Yorktown Heights, NY 10598

Dr. Dennis Kibler
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. David Kieras
University of Michigan
Technical Communication
College of Engineering
1223 E. Engineering Building
Ann Arbor, MI 48109

Dr. Peter Kincaid
Training Analysis
& Evaluation Group
Department of the Navy
Orlando, FL 32813

Dr. Walter Kintsch
Department of Psychology
University of Colorado
Campus Box 345
Boulder, CO 80302

Dr. David Klahr
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Mazie Knerr
Program Manager
Training Research Division
HumRRO
1100 S. Washington
Alexandria, VA 22314

Dr. Janet L. Kolodner
Georgia Institute of Technology
School of Information
& Computer Science
Atlanta, GA 30332

Dr. Stephen Kosslyn
Harvard University
1236 William James Hall
33 Kirkland St
Cambridge, MA 02138

Dr. Kenneth Kotovsky
Department of Psychology
Community College of
Allegheny County
800 Allegheny Avenue
Pittsburgh, PA 15233

Dr. David H. Krantz
2 Washington Square Village
Apt. # 15J
New York, NY 10012

Dr. Benjamin Kuipers
University of Texas at Austin
Department of Computer Sciences
T. S. Painter Hall 3.28
Austin, Texas 78712

Dr. David R. Lambert
Naval Ocean Systems Center
Code 441T
271 Catalina Boulevard
San Diego, CA 92152-6800

Distribution List [Illinois/Gentner] NR 667-551

Dr. Pat Langley
University of California
Department of Information
and Computer Science
Irvine, CA 92717

Dr. Clayton Lewis
University of Colorado
Department of Computer Science
Campus Box 430
Boulder, CO 80309

Dr. Barbara Means
Human Resources
Research Organization
1100 South Washington
Alexandria, VA 22314

Dr. Marcy Lansman
University of North Carolina
The L. L. Thurstone Lab.
Davis Hall 013A
Chapel Hill, NC 27514

Matt Lewis
Department of Psychology
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Douglas L. Medin
Department of Psychology
University of Illinois
603 E. Daniel Street
Champaign, IL 61820

Dr. Jill Larkin
Carnegie-Mellon University
Department of Psychology
Pittsburgh, PA 15213

Library
Naval Training Systems Center
Orlando, FL 32813

Dr. George A. Miller
Department of Psychology
Green Hall
Princeton University
Princeton, NJ 08540

Dr. Jean Lave
School of Social Sciences
University of California
Irvine, CA 92717

Dr. Jane Malin
Mail Code SR 111
NASA Johnson Space Center
Houston, TX 77058

Dr. William Montague
NPRDC Code 13
San Diego, CA 92152-6800

Dr. Robert Lawler
Information Sciences, FRL
GTE Laboratories, Inc.
40 Sylvan Road
Waltham, MA 02254

Dr. William L. Maloy
Chief of Naval Education
and Training
Naval Air Station
Pensacola, FL 32508

Dr. Allen Munro
Behavioral Technology
Laboratories - USC
1845 S. Elena Ave., 4th Floor
Redondo Beach, CA 90277

Dr. Alan M. Lesgold
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Sandra P. Marshall
Dept. of Psychology
San Diego State University
San Diego, CA 92182

Chair, Department of
Computer Science
U.S. Naval Academy
Annapolis, MD 21402

Dr. Jim Levin
Dept. of Educational Psy
210 Education Building
1310 South Sixth St.
Champaign, IL 61810-6990

Dr. Manton M. Matthews
Department of Computer Science
University of South Carolina
Columbia, SC 29208

Dr. Allen Newell
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. John Levine
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Dr. Richard E. Mayer
Department of Psychology
University of California
Santa Barbara, CA 93106

Dr. Richard E. Nisbett
University of Michigan
Institute for Social Research
Room 5261
Ann Arbor, MI 48109

Dr. Michael Levine
Educational Psychology
210 Education Bldg
University of Illinois
Champaign, IL 61820

Dr. Joe McLachlan
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. James McMichael
Assistant for MPT Research,
Development, and Studies
OP 01B7
Washington, DC 20370

Dr. Mary Jo Nissen
University of Minnesota
N218 Elliott Hall
Minneapolis, MN 55455

Director, Training Laboratory
NPRDC (Code 05)
San Diego, CA 92152-6800

Distribution List [Illinois/Gentner] NR 667-551

Director, Manpower and Personnel
Laboratory,
NPRDC (Code 06)
San Diego, CA 92152-6800

Special Assistant for Marine
Corps Matters,
ONR Code 00MC
800 N Quincy St.
Arlington, VA 22217-5000

Military Assistant for Training and
Personnel Technology,
OUSD (R & E)
Room 3D129, The Pentagon
Washington, DC 20301-3080

Director, Human Factors
& Organizational Systems Lab,
NPRDC (Code 07)
San Diego, CA 92152-6800

Psychologist
Office of Naval Research
Liaison Office, Far East
APO San Francisco, CA 96503

Dr. David N Perkins
Educational Technology Center
337 Gutman Library
Appian Way
Cambridge, MA 02138

Fleet Support Office,
NPRDC (Code 301)
San Diego, CA 92152-6800

Dr. Judith Orasanu
Army Research Institute
5001 Eisenhower Avenue
Alexandria, VA 22303

Dr. Nancy Perry
Chief of Naval Education
and Training, Code 00A2A
Naval Station Pensacola
Pensacola, FL 32508

Library, NPRDC
Code P201L
San Diego, CA 92152-6800

Prof. Seymour Papert
20C-109
Massachusetts Institute
of Technology
Cambridge, MA 02139

Department of Computer Science
Naval Postgraduate School
Monterey, CA 93940

Dr. Harold F. O'Neil, Jr.
School of Education - WPH 801
Department of Educational
Psychology & Technology
University of Southern California
Los Angeles, CA 90089-0031

Dr. James Paulson
Department of Psychology
Portland State University
P O Box 751
Portland, OR 97207

Dr. Steven Pinker
Department of Psychology
E10-018
MIT
Cambridge, MA 02139

Dr. Michael Oberlin
Naval Training Systems Center
Code 711
Orlando, FL 32813-7100

Dr. Roy Pea
Bank Street College of
Education
610 W 112th Street
New York, NY 10025

Dr. Tjeerd Plomp
Twente University of Technology
Department of Education
P O Box 217
7500 AE ENSCHEDE
THE NETHERLANDS

Dr. Stellan Ohlsson
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Dr. Douglas Pearse
DCIEM
Box 2000
Downsview, Ontario
CANADA

Dr. Martha Polson
Department of Psychology
Campus Box 346
University of Colorado
Boulder, CO 80309

Office of Naval Research,
Code 1133
800 N Quincy Street
Arlington, VA 22217-5000

Dr. James W. Pellegrino
University of California,
Santa Barbara
Department of Psychology
Santa Barbara, CA 93106

Dr. Peter Polson
University of Colorado
Department of Psychology
Boulder, CO 80309

Office of Naval Research,
Code 1142
800 N Quincy St.
Arlington, VA 22217-5000

Dr. Virginia E. Pendergrass
Code 711
Naval Training Systems Center
Orlando, FL 32813-7100

Dr. Steven E. Poltrock
MCC
9430 Research Blvd.
Echelon Bldg #1
Austin, TX 78759-6509

Psychologist
Office of Naval Research
Branch Office, London
Box J9
FPO New York, NY 09510

Distribution List [Illinois/Gentner] NR 667-551

Dr. Harry E. Pople
University of Pittsburgh
Decision Systems Laboratory
1360 Scaife Hall
Pittsburgh, PA 15261

Dr. Mary C. Potter
Department of Psychology
MIT (E-10-032)
Cambridge, MA 02139

Dr. Joseph Psotka
ATTN PERI-1C
Army Research Institute
5001 Eisenhower Ave.
Alexandria, VA 22333

Dr. Lynne Reder
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Dr. James A. Reggia
University of Maryland
School of Medicine
Department of Neurology
22 South Greene Street
Baltimore, MD 21201

Dr. Fred Reif
Physics Department
University of California
Berkeley, CA 94720

Dr. Lauren Resnick
Learning R & D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15213

Dr. Gil Ricard
Mail Stop C04-14
Grumman Aerospace Corp.
Bethpage, NY 11714

Mark Richer
1041 Lake Street
San Francisco, CA 94118

Dr. Mary S. Riley
Program in Cognitive Science
Center for Human Information
Processing
University of California
La Jolla, CA 92093

Dr. Linda G. Roberts
Science, Education, and
Transportation Program
Office of Technology Assessment
Congress of the United States
Washington, DC 20510

Dr. William B. Rouse
Search Technology, Inc.
25-b Technology Park/Atlanta
Norcross, GA 30092

Dr. David Rumelhart
Center for Human
Information Processing
Univ. of California
La Jolla, CA 92093

Dr. Roger Schank
Yale University
Computer Science Department
P O Box 2158
New Haven, CT 06520

Dr. Walter Schneider
Learning R&D Center
University of Pittsburgh
3939 O'Hara Street
Pittsburgh, PA 15260

Dr. Alan H. Schoenfeld
University of California
Department of Education
Berkeley, CA 94720

Dr. Janet Schobeld
Learning R&D Center
University of Pittsburgh
Pittsburgh, PA 15260

Karen A. Schriver
Department of English
Carnegie-Mellon University
Pittsburgh, PA 15213

Dr. Judah L. Schwartz
MIT
20C-120
Cambridge, MA 02139

Dr. Marc Sebrechts
Department of Psychology
Wesleyan University
Middletown, CT 06475

Dr. Judith Segal
OERI
555 New Jersey Ave., NW
Washington, DC 20208

Dr. Sylvia A. S. Shafto
Department of
Computer Science
Towson State University
Towson, MD 21204

Dr. Ben Shneiderman
Dept. of Computer Science
University of Maryland
College Park, MD 20742

Dr. Lee Shulman
Stanford University
1040 Cathcart Way
Stanford, CA 94305

Dr. Robert S. Siegler
Carnegie-Mellon University
Department of Psychology
Schenley Park
Pittsburgh, PA 15213

Dr. Derek Sleeman
Stanford University
School of Education
Stanford, CA 94305

Dr. Edward E. Smith
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Richard E. Snow
Department of Psychology
Stanford University
Stanford, CA 94306

Distribution List [Illinois/Gentner] NR 007-551

Dr. Elliot Soloway
Yale University
Computer Science Department
P O Box 2158
New Haven, CT 06520

Chair, Department of
Computer Science
Towson State University
Towson, MD 21204

Dr. Michael Williams
IntelliCorp
1975 El Camino Real West
Mountain View, CA 94040-2216

Dr. Richard Sorensen
Navy Personnel R&D Center
San Diego, CA 92152-6800

Chair, Department of
Psychology
Towson State University
Towson, MD 21204

Dr. Robert A. Wisher
U.S. Army Institute for the
Behavioral and Social Sciences
5001 Eisenhower Avenue
Alexandria, VA 22333

Dr. Kathryn T. Spoehr
Brown University
Department of Psychology
Providence, RI 02912

Dr. Kurt Van Lehn
Department of Psychology
Carnegie-Mellon University
Schenley Park
Pittsburgh, PA 15213

Mr. John H. Wolfe
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Robert Sternberg
Department of Psychology
Yale University
Box 11A, Yale Station
New Haven, CT 06520

Dr. Beth Warren
Bolt Beranek & Newman, Inc.
50 Moulton Street
Cambridge, MA 02138

Dr. Wallace Wulfeck, III
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Albert Stevens
Bolt Beranek & Newman, Inc.
10 Moulton St
Cambridge, MA 02238

Dr. Donald Weitsman
MITRE
1820 Dolley Madison Blvd.
MacLean, VA 22102

Dr. Joe Yasatuke
AFHRL/LRT
Lowry AFB, CO 80230

Dr. Thomas Sticht
Navy Personnel R&D Center
San Diego, CA 92152-6800

Dr. Keith T. Wescourt
FMC Corporation
Central Engineering Labs
1185 Coleman Ave., Box 580
Santa Clara, CA 95052

Dr. Masoud Yazdani
Dept. of Computer Science
University of Exeter
Exeter EX4 4QL
Devon, ENGLAND

Dr. John Tangney
AFOSR/NL
Boiling AFB, DC 20332

Dr. Douglas Wetzel
Code 12
Navy Personnel R&D Center
San Diego, CA 92152-6800

Mr. Carl York
System Development Foundation
181 Lytton Avenue
Suite 210
Palo Alto, CA 94301

Dr. Kikumi Tatsuoka
CERL
252 Engineering Research
Laboratory
Urbana, IL 61801

Dr. Barbara White
Bolt Beranek & Newman, Inc.
10 Moulton Street
Cambridge, MA 02238

Dr. Joseph L. Young
Memory & Cognitive
Processes
National Science Foundation
Washington, DC 20550

Dr. Perry W. Thorndyke
FMC Corporation
Central Engineering Labs
1185 Coleman Avenue, Box 580
Santa Clara, CA 95052

Dr. Christopher Wickens
Department of Psychology
University of Illinois
Champaign, IL 61820

Dr. Douglas Towne
Behavioral Technology Labs
1845 S. Elena Ave
Redondo Beach, CA 90277

Dr. Heather Wild
Naval Air Development Center
Code 6021
Warminster, PA 18974-5000

END

2-87

DTIC